

Topological model selection: a case-study in tumour-induced angiogenesis

Robert A McDonald,¹ Helen M Byrne,^{1,2} Heather A Harrington,^{1,3,4,5}
Thomas Thorne^{6,*} and Bernadette J Stolz^{7,*}

¹Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter, Woodstock Road, Oxford, OX2 6GG, United Kingdom, ²Ludwig Institute for Cancer Research, Nuffield Department of Medicine, Old Road Campus Research Building, Oxford OX3 7DQ, United Kingdom, ³Faculty of Mathematics, Technische Universität Dresden, 01062 Dresden, Germany, ⁴Centre for Systems Biology Dresden (CSBD), Pfotenhauerstrasse 108, 01062 Dresden, Germany, ⁵Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG), 01307 Dresden, Germany, ⁶Computer Science Research Centre, University of Surrey, GU2 7XH, United Kingdom and ⁷Department of Machine Learning and Systems Biology, Max Planck Institute of Biochemistry, Am Klopferspitz 18, 82152 Martinsried, Germany

*Corresponding authors. tom.thorne@surrey.ac.uk, stolz@biochem.mpg.de

Abstract

Motivation: Comparing mathematical models offers a means to evaluate competing scientific theories. However, exact methods of model calibration are not applicable to many probabilistic models which simulate high-dimensional spatio-temporal data. Approximate Bayesian Computation is a widely-used method for parameter inference and model selection in such scenarios, and it may be combined with Topological Data Analysis to study models which simulate data with fine spatial structure.

Results: We develop a flexible pipeline for parameter inference and model selection in spatio-temporal models. Our pipeline identifies topological summary statistics which quantify spatio-temporal data and uses them to approximate parameter and model posterior distributions. We validate our pipeline on models of tumour-induced angiogenesis, inferring four parameters in three established models and identifying the correct model in synthetic test-cases.

Availability and implementation: Simulation code for all models, data analyses, parameter inference and model selection is available online at <https://github.com/rmcdmaths/tms/>.

Supplementary information: *Supplementary Information* will be available online.

Key words: spatial modelling, topology, parameter inference, model selection

1. Introduction

Given multiple mathematical models which aim to reproduce the same biological data, determining which model and parameters give the best fit presents a theoretical and computational challenge. For example, spatio-temporal models often simulate complex high-dimensional data which is difficult to quantify and compare to observed data. Such models do not in general yield tractable likelihood functions, which significantly hinders the use of exact methods for parameter inference and model selection (Kirk *et al.*, 2013).

Many mathematical models have been developed to study the mechanisms underlying tumour-induced angiogenesis (Scianna *et al.*, 2013; Vilanova *et al.*, 2017), a hallmark of cancer (Hanahan and Weinberg, 2011). Tumour cells use chemical

signals to stimulate the growth of new blood vessels from existing vasculature (Ferrara, 2002), which provide a tumour mass with oxygen and nutrients that it requires to grow. However, instead of concise equations determining the growth of angiogenic networks, such models often comprise multiple agents and heterogeneous environments whose interactions depend non-deterministically on their spatial organisation. Discrete models of tumour-induced angiogenesis, for example, use multiple model rules and parameters to reproduce the branches, loops, and multiple components that characterise real vascular networks.

We use Topological Data Analysis (TDA), Approximate Bayesian Computation (ABC), and Random Forests (RFs) to develop a pipeline for parameter inference and model selection applicable to spatio-temporal models. TDA offers a toolkit of methods for quantifying spatial data and has previously been

used to study multi-agent temporal systems (Topaz et al., 2015; Bhaskar et al., 2019; Stolz et al., 2024). In related work, TDA has been used to compare models of insect locomotion (Ulmer et al., 2019) and pattern formation in zebrafish (Cleveland et al., 2023). ABC provides a statistical framework for using model simulations to approximate posterior distributions when likelihood functions are not available (Lintusaari et al., 2016). RFs are an ensemble estimation method from machine learning which have previously been combined with ABC to estimate parameter values (Raynal et al., 2018) and rank candidate models (Pudlo et al., 2015).

We begin by outlining three existing models of tumour-induced angiogenesis in which exact methods of parameter inference and model selection are not applicable. We show how TDA can be used to characterise spatial data simulated by the models and briefly describe methods from ABC and RF. We then present a three-step pipeline for parameter inference and model selection which we apply to the angiogenesis models. Commenting on the applicability of our pipeline to experimental data, we discuss how topological summaries may be used to evaluate a variety of modelling approaches in biology.

2. Model data and analysis

2.1. Angiogenesis models

Discrete models of tumour-induced angiogenesis simulate the movement of individual Endothelial Cells (ECs). Vascular Endothelial Growth Factors (VEGF) produced by tumour cells initiate a cascade of chemical reactions which drive ECs towards the tumour. Early models employed the *snail-trail model* (Balding and McElwain, 1985) in which tip ECs migrate up spatial gradients of VEGF and fibronectin, while stalk ECs proliferate in their path to produce a contiguous line of cells. When two separate trails of ECs meet they fuse together to form a loop, in a process known as anastomosis. A tip EC may also split into two tip ECs, which thereafter move independently. Migrating, branching and looping tip ECs eventually reach the tumour mass and the connecting trails of stalk EC form a blood vessel network. Recent models reflect modern discoveries of cell mixing and phenotype switching (Stepanova et al., 2024), where ECs change type and overtake each other before forming a stable blood vessel network. Other models view ECs as a continuous population density rather than individual cells (Martinson et al., 2021) and account for blood flow and nutrient delivery when simulating vascular networks (Hormuth et al., 2021). We develop our pipeline of parameter inference and model selection on discrete angiogenesis models due to their simulation of finely resolved spatial data.

The Anderson-Chaplain (AC) (Anderson and Chaplain, 1998), Stokes-Lauffenberger (SL) (Stokes et al., 1991), and Plank-Sleeman (PS) (Plank and Sleeman, 2004) models employ the snail-trail model to simulate movement of individual ECs in a two-dimensional, square domain. We assume that VEGF levels increase from the bottom of the domain to a tumour at the top, guiding ECs to move upwards. Each model initialises multiple distinct tip ECs along the bottom of the domain, simulating their trajectories according to model-specific movement rules. Figure 1 illustrates the movement rules and model parameters in each model. See *Supplementary Information* Section A for full statements of each model and their parameters.

2.2. Data Generation and Analysis

Each angiogenesis model described in Section 2.1 simulates EC movement in a square domain, and we wish to quantify the spatial properties of each simulation. To do so, we overlay a regular grid onto the simulation domain at the final simulation timestep and compute a collection of statistics on the point cloud consisting of the grid locations of ECs. We compute the mean, standard deviation, minimum, maximum, range, interquartile range, and the 10th, 25th, 75th and 90th percentiles of the x and y co-ordinates of the EC grid locations. These *spatially-averaged* statistics measure how simulated ECs are distributed along the vertical and horizontal axes of the simulation domain, and were previously used for comparison to topological statistics by Nardini et al. (2021) to distinguish the AC model’s behaviour in different parameter regimes. To quantify spatial structures in simulated vasculature which may be obscured by spatially-averaged statistics (such as loops, branches and separated components) we use Extended Persistent Homology (EPH).

EPH arises from Persistent Homology (PH), a common pipeline within TDA for quantifying spatial features in data. To compute PH, spatial data must be converted into a sequence of simplicial complexes known as a filtration (see, for example, Otter et al. (2017)). Intuitively, a simplicial complex is a graph that includes nodes, edges, and higher-order connections such as triangles or tetrahedra. A filtration is a sequence of embedded simplicial complexes Σ_k such that $\Sigma_k \subseteq \Sigma_{k+1}$ for $k = 0, 1, \dots, K-1$. The dimension- p PH of a filtration is the sequence (1) of \mathbb{F} -vector spaces $H_p(\Sigma_k)$ referred to as homology groups. Throughout this analysis, we work with the field $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$, which is widely adopted in applications for its simplicity and interpretability. The homology group $H_p(\Sigma_k)$ captures topological features of dimension p in Σ_k . Intuitively, H_0 detects connected components, H_1 detects loops and, in general, H_p detects p -dimensional holes (Carlsson, 2009). The inclusion maps between successive simplicial complexes in the filtration induce linear maps in sequence (1) which allow us to track p -dimensional holes across the filtration.

$$H_p(\Sigma_0) \rightarrow \dots \rightarrow H_p(\Sigma_k) \rightarrow \dots \rightarrow H_p(\Sigma_K) \quad (1)$$

Sequence (1) uniquely decomposes which allows us to represent each topological feature as a persistence pair (b, d) (Zomorodian and Carlsson, 2005). The birth b corresponds to the filtration index k in the sequence at which a p -dimensional hole is first created in the simplicial complex, and the death d is the filtration index in which this hole is filled in. The persistence $d - b$ measures how long the corresponding feature persists through the filtration. Note that persistence may be infinite if the p -dimensional hole is never filled in (and $d = \infty$). The interpretation of (b, d) depends on how the filtration was constructed from the underlying spatial data. For example, Nardini et al. (2021) used a *sweeping-plane* filtration to analyse the AC model, in which persistence pairs quantified the location of loops and components in simulated vascular networks (measured in distance from the tumour).

Instead of using a standard PH pipeline, we compute Extended Persistent Homology (EPH) (Cohen-Steiner et al., 2009). Like PH, EPH uses a filtration to compute the sequence (1). In contrast to PH, EPH uses additional simplicial complexes to extend the sequence of vector spaces providing refined topological descriptors of spatial data. Specifically, EPH uses a sequence of simplicial complexes Σ^k containing those simplices that are in Σ_K but not in Σ_k . The relative homology $H_p(\Sigma_K, \Sigma^k)$ detects topological

features in the quotient complex Σ_K/Σ^k (Edelsbrunner and Harer, 2022), and since $\Sigma^k \supseteq \Sigma^{k+1}$ for $k = K-1, \dots, 1, 0$, quotient maps between these complexes induce linear maps tracking topological features through sequence (2).

$$H_p(\Sigma_K, \Sigma^K) \rightarrow \dots \rightarrow H_p(\Sigma_K, \Sigma^k) \rightarrow \dots \rightarrow H_p(\Sigma_K, \Sigma^0) \quad (2)$$

The concatenation of sequences (1) and (2) is known as the dimension- p EPH of the filtration. p -dimensional holes may be detected and then filled within either sequence (1) or (2) and the corresponding indices k may be extracted as persistence pairs (b, d) . We compute dimension $p = 0$ and $p = 1$ EPH of two sweeping-plane filtrations to quantify spatial features in data simulated by the AC, SL or PS models. Using these filtrations, extended persistence pairs quantify the size and location of components and loops in simulated vasculature, in addition to finding branches which emerge from vessel components when ECs bifurcate. Each sweeping-plane filtration measures these topological features in either the horizontal (x) or vertical (y) co-ordinate direction. See *Supplementary Information* Section B for details of the filtrations we use to compute EPH and an example computation on an angiogenesis simulation. We convert extended persistence pairs in each dimension into fixed-length vectors using two common vectorizations. Persistence Images (Adams et al., 2017) assign a weight to each extended persistence pair and create a smooth surface which is discretised over a grid to produce a stable, fixed-length vector. Persistence statistics (Ali et al., 2023) are vectors that consist of the mean, standard deviation, median, interquartile range, full range, and the 10th, 25th, 75th and 90th percentiles computed from the births b , deaths d , persistences $d-b$ and midpoints $(b+d)/2$ of persistence pairs (b, d) .

3. Methods

3.1. Approximate Bayesian Computation

ABC provides a statistical framework for using data to infer model parameters. Suppose a model uses parameters Θ to simulate data \mathcal{D} according to some probability distribution $p(\mathcal{D}|\Theta)$, called the *likelihood*. Parameter inference aims to determine the *posterior* distribution $p(\Theta|\mathcal{D})$, which is the probability that parameters Θ generated observed data \mathcal{D} . Using previous experiments or assumptions about feasible parameter values, one may define a *prior* distribution $p(\Theta)$ representing knowledge of the parameter values before data has been observed. The likelihood, prior and posterior are related by Bayes' rule,

$$p(\Theta|\mathcal{D}) = \frac{p(\mathcal{D}|\Theta)p(\Theta)}{p(\mathcal{D})},$$

where the evidence $p(\mathcal{D})$ is the integral $\int_{\Theta} p(\mathcal{D}|\Theta)p(\Theta)d\Theta$ over parameters Θ in the support of the prior. Although Bayes' rule gives a closed formula for the posterior distribution, it is often impractical to use directly. The likelihood function $p(\mathcal{D}|\Theta)$ may be too complicated to derive for probabilistic spatial models in which many datasets \mathcal{D} may be simulated from the same parameters Θ . Instead, Bayes' rule is used to derive Algorithm 1 which allows sampling from the posterior, even when the likelihood and evidence are not known. When the tolerance ϵ in Algorithm 1 is set to zero, the distribution of accepted parameters is (a constant multiple of) the posterior distribution $p(\Theta|\mathcal{D}^*)$ (Frazier et al., 2018). However, it is often not appropriate or possible to seek

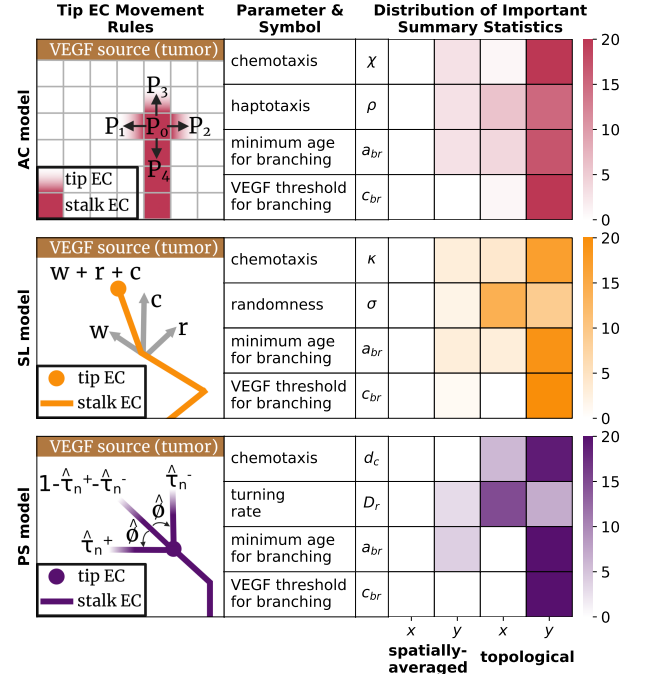


Fig. 1: In the AC model, a tip EC makes one of five possible moves on a square lattice in each time-step according to probabilities P_0, P_1, P_2, P_3, P_4 . A chemotaxis parameter biases movement probabilities in the direction of increasing VEGF concentration, and a haptotaxis parameter biases moves in the direction of increasing fibronectin. In the SL model, tip ECs move in any direction (off-lattice) with velocities modelled by a two-dimensional Stochastic Differential Equation (SDE). Parameters κ and σ determine how strongly an EC's current velocity w is affected by the VEGF gradient c and random variation r respectively. The PS model assigns a constant speed to each EC and, at each time-step, rotates the angle that the velocity vector makes with the vertical by $\hat{\phi}$. The probability $\hat{\tau}_n^+ + \hat{\tau}_n^-$ of turning is determined by a turning rate parameter, and a chemotaxis parameter biases turns that re-orient the EC's direction towards the tumour. In all models, a tip EC may bifurcate into two ECs which thereafter move independently if its age exceeds the minimum age for branching parameter a_{br} and the VEGF concentration at its location exceeds the the VEGF threshold for branching parameter c_{br} . We show how many spatially-averaged and topological statistics, computed in either the x or y co-ordinate direction, appear among the 100 most important summary statistics to the inference of each parameter.

an exact posterior distribution from observed data, since a model may rarely reproduce observed data \mathcal{D}^* exactly, and the observed data may be noisy. It is therefore advisable to choose ν and ϵ such that parameter values are accepted if they simulate data that is similar to observed data. The general form of such a distance function is $\nu(\mathcal{D}^*, \mathcal{D}_i) = \|X^* - X_i\|_2$ where X_i and X^* are vectors of *summary statistics* computed from model data \mathcal{D}_i and observed data \mathcal{D}^* respectively (this distance may be averaged over several instances of observed data \mathcal{D}^*). Summary statistics aim to capture relevant properties of data as a low dimensional vector. As the tolerance ϵ approaches 0, the distribution of parameters accepted

by Algorithm 1 approaches $p(\Theta|X^*)$, which equals $p(\Theta|\mathcal{D}^*)$ if the summary statistics are sufficient for the model in question, or is a close approximation if the summary statistics are insufficient but informative (Joyce and Marjoram, 2008). However, if ϵ is too small, few parameters will be accepted and many simulations may be needed to approach the true posterior. To overcome the difficulty of choosing a single tolerance, ABC with Sequential Monte Carlo sampling (ABC-SMC) modifies Algorithm 1, and instead generates a sequence of immediate distributions which approach the true posterior using tolerances $\epsilon_0 > \epsilon_1 > \dots > \epsilon_{\text{final}} > 0$. For example, the ABC-SMC algorithm of Del Moral et al. (2012) generates the schedule of tolerances by decreasing the proportion of accepted parameters in each intermediate distribution by a fixed factor $\alpha \in (0, 1)$. We discuss the choice of summary statistics and ν in Section 4.

Algorithm 1 ABC with rejection sampling

Input: Observed data \mathcal{D}^* , a model which generates data \mathcal{D} from parameters Θ with prior $p(\Theta)$, a distance function $\nu : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ comparing simulated and observed data, and a small tolerance $\epsilon > 0$.

Output: A collection of samples θ_i from the posterior $p(\Theta|\mathcal{D}^*)$

- 1: **for** Candidate parameter values θ_i sampled from $p(\Theta)$ **do**
- 2: Simulate data \mathcal{D}_i from the model using parameter value θ_i
- 3: **if** $\nu(\mathcal{D}^*, \mathcal{D}_i) \leq \epsilon$ **then**
- 4: accept θ_i
- 5: **end if**
- 6: **end for**

3.2. Random Forests

Random Forests (RFs) (Breiman, 2001) learn relationships between feature vectors and response variables. Training data comprising a collection of feature vectors $X_i \in \mathcal{X}$ and corresponding response variables $y_i \in \mathcal{Y}$ are used to train a RF, enabling it to predict the true response variable y^* of an unseen feature vector X^* . Regression RFs are used when y_i are continuous values and classification RFs are used when y_i are discrete labels. Raynal et al. (2018) used a regression RF for parameter inference by using simulated data \mathcal{D}_i to learn the relationship between summary statistics X_i and parameter values $y_i = \theta_i$. Given summary statistics X^* of unseen data \mathcal{D}^* , the prediction $RF(X^*)$ predicts the true parameter value θ^* . In addition to predicting unseen feature vectors, a trained RF provides useful information about the training data. The out-of-bag prediction $RF_{\text{ob}}(X_i)$ estimates the (known) response variable y_i using pairs from the training data other than (X_i, y_i) . The out-of-bag error rate $p(RF(X_i) \neq y_i)$ then gives a (unbiased) measure of how well the relationship between X_i and y_i is captured by the rest of the training data. A trained RF also gives a measure of the *importance* of each co-variate j within feature vectors $X_i = (X_i^0, \dots, X_i^j, \dots, X_i^{n_t})$ to the problem of predicting response variable y_i . Intuitively, important features are those whose values within X_i and X'_i differ when y_i and y'_i do, and which are hence useful in learning the relationship between training data \mathcal{X} and \mathcal{Y} .

3.3. Model Selection

Given observed data \mathcal{D}^* , the model posterior $p(m|\mathcal{D}^*)$ is the probability that models $m = m_i$ generated \mathcal{D}^* for $m_i =$

$1, 2, \dots$. Algorithm 1 and its extensions rely on the approximation $p(\Theta|X^*) \approx p(\Theta|\mathcal{D}^*)$, which holds as long as the vector X_i carries a similar amount of information about the parameter value θ_i as the simulated data \mathcal{D}_i itself. However, the information loss suffered by a collection of summary statistics may vary between models (Robert et al., 2011), so it is inadvisable simply to infer m_i as a (discrete) parameter using Algorithm 1. Pudlo et al. (2015) instead used two RFs to approximate $p(m|\mathcal{D}^*)$. A classification RF learns the relationship between simulated data X_i and model label $y_i = m_i$ and gives a prediction $RF(X^*)$ of the model m^* which generated unseen data \mathcal{D}^* . A regression RF is then trained to learn the relationship between X_i and $p(RF_{\text{ob}}(X_i) \neq m_i)$ —the out-of-bag error rate of the classification RF. The regression RF is then used to estimate posterior probability $p(m = m^*|\mathcal{D}^*)$ as $1 - p(RF(X^*) \neq m^*)$.

4. Spatial Parameter Inference and Model Selection

Given observed data \mathcal{D}^* , we wish to approximate the parameter posterior $p(\Theta|\mathcal{D}^*)$ for candidate models $m = m_i$ for $i = 1, 2, \dots$ and the model posterior $p(m|\mathcal{D}^*)$. Informative summary statistics may be used to define a distance function $\nu(\mathcal{D}^*, \mathcal{D}_i)$ to infer parameters using ABC. However, uninformative or poorly scaled summary statistics may misrepresent the difference between datasets generated by similar parameters (Blum et al., 2013). We therefore seek a collection of summary statistics that quantify simulated data and, in particular, quantify how simulated data changes when different model rules and parameters are used to generate it. We use informative summary statistics to approximate parameter and model posteriors in a three-step pipeline. We test this pipeline on toy models in *Supplementary Information* Section D and apply it to the three angiogenesis models in Section 5.

Step 1: Identify informative summary statistics for each model

We draw parameter values θ_i from the prior distribution $p(\Theta)$ for each parameter in each model, simulate model data \mathcal{D}_i , and compute spatially-averaged and topological statistics X_i from the final simulated time-step. We train regression RFs to learn the relationship between summary statistics X_i and parameter values θ_i —one RF for each parameter in each model. We then rank the spatially-averaged and topological statistics by their importance according to the RF (see Section 3.2). In each RF, feature importance decreases exponentially (as in Raynal et al. (2018)) and a small subset of summary statistics provides most of the predictive power of each RF. We select an equal number of informative summary statistics from each RF and collect a total of n_s for each model. We find that $n_s = 100$ is a good choice for the AC, SL and PS angiogenesis models. See *Supplementary Information* Section C for a full definition of RF feature importance and a discussion of how we choose n_s .

Step 2: Fit each model to the observed data

We use the distance function $\nu(\mathcal{D}^*, \mathcal{D}_i) = \|x^* - x_i\|_2$ to compare observed data \mathcal{D}^* to simulated data \mathcal{D}_i . Here, x_i is the vector of summary statistics X_i restricted to the top n_s features for each model, x^* is computed from observed data, and the distance is averaged over multiple instances of observed data. We scale each summary statistic by the largest absolute value of that summary statistic in the training data. Doing so ensures that each summary statistic contributes approximately equally to the

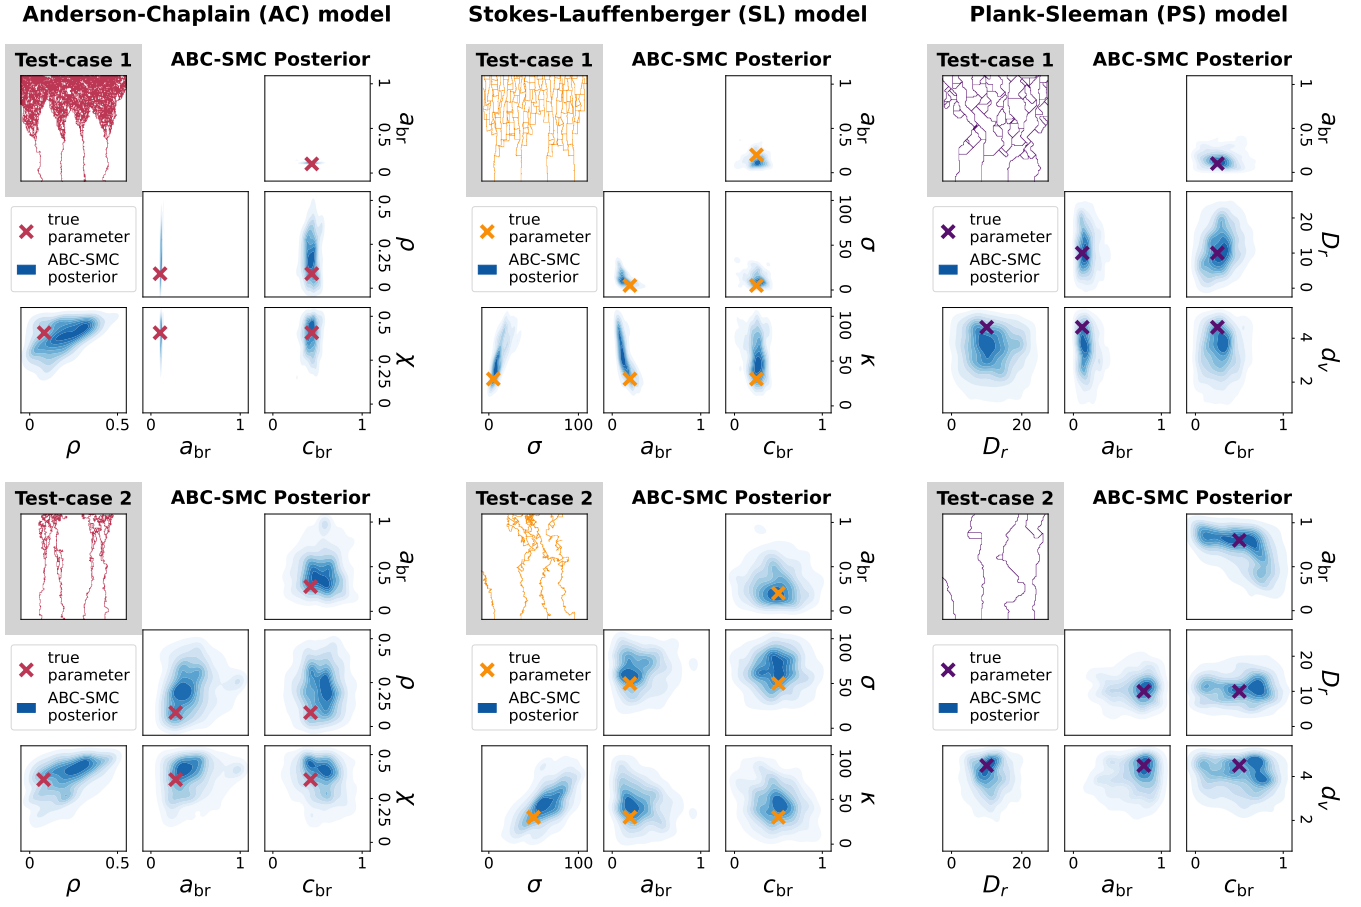


Fig. 2: We infer the minimum age for branching (a_{br}) and VEGF threshold for branching (c_{br}) in each model, as well as chemotaxis and haptotaxis parameters (χ and ρ) in the AC model, chemotaxis and randomness parameters (κ and σ) in the SL model, and chemotaxis and turning rate parameters (d_c and D_r) in the PS model. We simulate each model 10 times at known parameter values to generate two synthetic test-cases for each model, and show the final time-step of one such simulation. We then use Steps 1-2 of Section 3 to approximate the parameter posterior $p(\Theta|\mathcal{D}^*)$ in each test-case. We project the approximate ABC-SMC posterior to each parameter pair and plot the resulting distributions (fitting a Gaussian kernel to the parameter values accepted in the final population of the ABC-SMC algorithm), along with the true parameter which generated the test-case.

distance function ν and limits the influence of poorly scaled summary statistics. We then use the ABC-SMC algorithm of Del Moral *et al.* (2012) to approximate $p(\Theta|\mathcal{D}^*)$ for each model.

Step 3: Approximate the model posterior

Following Pudlo *et al.* (2015), we train a classification RF to learn the relationship between (unscaled) summary statistics X_i and model indices m_i in the training data. We modify X_i to contain only those summary statistics which appear among the n_s most important summary statistics for all models. We then train a regression RF to learn the relationship between X_i and $p(RF_{\text{obs}}(X_i) \neq m_i)$ —the probability that the predicted model index is incorrect. The classification RF gives an estimate $RF(X^*)$ of the model m^* that generated the observed data \mathcal{D}^* , and the regression RF is used to estimate $p(m = m^*|\mathcal{D}^*)$.

5. Results

5.1. RFs find small subsets of informative summary statistics

We sample $n = 10,000$ model parameters from uniform priors with ranges taken from existing literature, or by analysing each model's data generation rules (See *Supplementary Information* Section A for details). We select $n_s = 100$ summary statistics for each model and report the type of summary statistics selected—spatially-averaged or topological, computed in the horizontal or vertical direction—in Figure 1. A mixture of spatially-averaged and topological statistics are selected for each parameter, however there is a clear preference for topological statistics. More summary statistics computed in the vertical direction are selected than those computed in the horizontal direction, which is unsurprising, since most parameters moderate the movement of ECs upwards. Two exceptions are the randomness parameter σ of the SL model and the turning coefficient D_r in the PS model. These parameters cause ECs to deviate from their upward trajectory, and their RFs rely on summary statistics computed in the horizontal direction for inference.

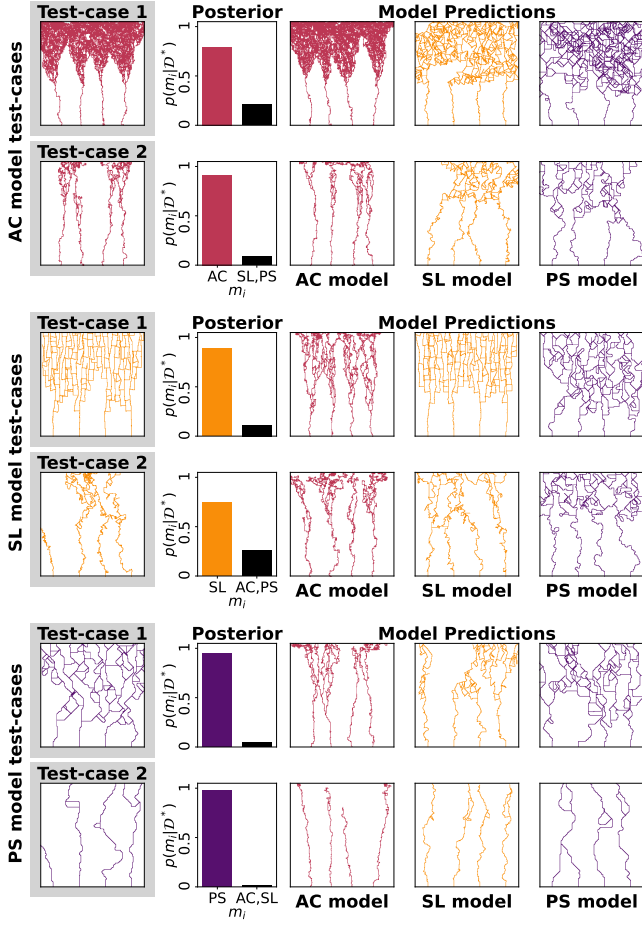


Fig. 3: We approximate the model posterior $p(m|\mathcal{D}^*)$ using the same four test-cases from Figure 2 (again highlighting one example of observed data \mathcal{D}^*). For each model, we show one example of simulated data associated to an inferred parameter in that model’s approximate parameter posterior. Each ‘prediction’ shows an example of that model’s approximation of the true data generation process.

5.2. ABC-SMC infers four parameters for each model and reproduces observed data

We create two test-cases for each model by simulating data at known parameter values 10 times. Each test-case exhibits quantitatively different vascular networks that each model can produce. We use the 100 summary statistics reported in Figure 1 to define a distance function between simulated and observed data as described in Section 3.1, averaging this distance across the 10 instances of observed data. We use ABC-SMC to infer four parameters for each model and show the resulting approximate posteriors in Figure 2. In each test-case, the approximate posterior is unimodal and encompasses the true parameter, often close to its densest part.

5.3. Random Forests correctly select models

Using the same test-cases, we perform Step 3, training a classification RF to predict the model index of unseen test data \mathcal{D}^* and a regression RF to predict the probability of mis-classifying

each test-case. Figure 3 shows the resulting approximate model posteriors, which identify the correct model with high probability in each test-case. We simulate each model at a parameter value drawn from its approximate parameter posterior and show the resulting ‘prediction’, which is a sample from the model’s approximation of the true data generating process. Each model generates data that is visually similar to the observed data, however we can identify the true model in each test-case.

6. Discussion

Using Topological Data Analysis (TDA) and Approximate Bayesian Computation (ABC), we have developed a pipeline for parameter inference and model selection applicable to complex spatial models. In previous work, TDA characterised the effect of two parameters in the AC model (Nardini et al., 2021) and was used along with ABC to infer them (Thorne et al., 2022). We extend this work by identifying a subset of informative summary statistics from multiple topological filtrations and using them to infer four parameters in three angiogenesis models using ABC-SMC. We further show how RFs can be used with TDA to approximate model posteriors and compare candidate models.

While we validated our pipeline using synthetic data simulated from angiogenesis models, several previous studies used *in vitro* data to inform model rules and parameters (Milde et al., 2013; Connor et al., 2015; Vergroesen et al., 2024). ABC provides a robust statistical framework for learning parameters and evaluating models, and TDA provides a variety of filtrations and vectorisations which may be adapted to different spatial data. The present work would therefore enhance previous model analysis, and in future work we will apply our methods to experimental data.

In this study, we considered three models in which the paths traced by tip EC form a static blood vessel network. In reality, sustained proliferation and vessel remodelling, where vasculature continually evolves after it is laid down, is characteristic of tumour-induced angiogenesis (Farnsworth et al., 2014). Indeed, vascular renormalisation, in which vessel-targeting agents prune small or inefficient blood vessels, is a theorised treatment strategy (Magnussen and Mills, 2021) aiming to temporarily enhance perfusion of the tumour to increase the effectiveness of radiotherapy (Köry et al., 2024). Topological invariants taking into account time-evolving data (Carlsson and de Silva, 2010), directed flow networks (Chaplin et al., 2024) and multiparameter filtrations (Vipond, 2020) could quantify such structural changes over time and be used to calibrate more sophisticated models.

Although we specialised our pipeline to discrete models of tumour-induced angiogenesis, its flexibility allows application to a range of spatio-temporal models. Any summary statistic which aims to quantify the desired properties of simulated data would be identified by the RF in Step 1 if it captures the effect of changing model parameters. In future, we will use our pipeline to robustly compare continuum, cell-based, agent-based, and discrete models by their ability to reproduce observed data.

A. Angiogenesis Models

We use the Anderson-Chaplain (AC) model (Anderson and Chaplain, 1998), the Stokes-Lauffenberger (SL) model (Stokes *et al.*, 1991), and the Plank-Sleeman (PS) model (Plank and Sleeman, 2004) as our case-study for topological model selection. In this section, we present the models in dimensionless form, giving formulae and simulation details for each model. We list model parameters in Table 1 and provide illustrative schematics in Figure 4.

Each model simulates the movement of multiple tip endothelial cells (ECs) in a two-dimensional square domain $\mathcal{I} = \{(x, y) : x, y \in [0, 1]\} \subset \mathbb{R}^2$. The position of a tip EC at time t is described by a variable $s^t = (s_x^t, s_y^t) \in \mathcal{I}$. We initialise four tip ECs along the lower edge of the domain (furthest from the tumour). We create four variables s such that $s^0 = (-1/8 + v/4, 0)$ for $v = 1, 2, 3, 4$. At discrete time intervals Δt , tip ECs migrate to a new position $s^{t+\Delta t} \in \mathcal{I}$, which is determined by a set of model-specific movement rules. In each model, we assume that a tumour is located along the domain's upper boundary ($y = 1$), acting as a source of vascular endothelial growth factors (VEGF). For simplicity, we prescribe the initial profile of VEGF $c(x, y, t = 0) = y$ in which the VEGF concentration decreases linearly with distance from the tumour. When a tip EC moves from s^t and $s^{t+\Delta t}$, the line segment connecting subsequent positions is assumed to be occupied thereafter by immobile stalk EC (*the snail-trail model* (Balding and McElwain, 1985)). Each model simulates EC movement over the (dimensionless) time interval $\mathcal{T} = [0, t_{\text{final}}]$. We set $\Delta t = 0.01$ and $t_{\text{final}} = 4$ meaning that model simulations terminate after 400 time-steps.

A.1. Anderson-Chaplain (AC) model

The AC model employs an on-lattice biased random walk to simulate tip EC movement in response to local levels of VEGF and fibronectin. A regular grid of $N \times N$ points spaced at intervals of h is placed on the square domain \mathcal{I} (we fix $N = 200$ and $h = 0.05$), and tip ECs move through the domain one grid space at a time. The initial concentration of VEGF is as defined above and the initial concentration of fibronectin is $f(x, y, 0) = 1 - y$. The time evolution of the VEGF and fibronectin concentrations, and the related probability that a tip EC moves left, right, down or up on the lattice, are derived from a system of Partial Differential Equations (PDEs).

$$\frac{\partial e}{\partial t} = D\Delta e - \chi \nabla \cdot (e \nabla c) - \rho \nabla \cdot (e \nabla f) \quad (3)$$

$$\frac{\partial f}{\partial t} = \beta e - \gamma e f \quad (4)$$

$$\frac{\partial c}{\partial t} = -\eta e c \quad (5)$$

The PDEs (3)–(5) describe how the spatial distribution of ECs, VEGF and fibronectin ($e, c, f : \mathcal{I} \times \mathcal{T} \rightarrow \mathbb{R}$) evolve over time. The parameter D determines the rate of EC random motility/diffusion, and parameters χ and ρ give the strength of the ECs' chemotactic and haptotactic responses to spatial gradients of VEGF and fibronectin respectively. ECs produce and degrade fibronectin at rates β and γ , and consume VEGF at rate η . The PDEs are closed by imposing no-flux boundary conditions along each side of the square domain \mathcal{I} .

To generate update rules for the concentration of VEGF and fibronectin at lattice points, as well as movement rules for the tip ECs, the PDE system is discretised using the Euler finite difference approximation. Let $e_{l,m}^t, f_{l,m}^t, c_{l,m}^t$ be the values of e, f and c at lattice points $(lh, mh) \in \mathcal{I}$ at time t . Discretising Equations (3)–(5) gives:

$$e_{l,m}^{t+\Delta t} = e_{l,m}^t P_0 + e_{l+1,m}^t P_1 + e_{l-1,m}^t P_2 + e_{l,m+1}^t P_3 + e_{l,m-1}^t P_4, \quad (6)$$

$$f_{l,m}^{t+\Delta t} = f_{l,m}^t (1 - \Delta t \gamma e_{l,m}^t) + \Delta t \beta e_{l,m}^t, \quad (7)$$

$$c_{l,m}^{t+\Delta t} = c_{l,m}^t (1 - \Delta t \eta e_{l,m}^t). \quad (8)$$

On the lattice, the VEGF and fibronectin initial conditions become $c_{l,m}^0 = m/N$ and $f_{l,m}^0 = 1 - m/N$. The discretisations (7)–(8) may then be used to update the VEGF and fibronectin concentrations $c_{l,m}^t$ and $f_{l,m}^t$ at grid locations (l, m) in discrete time-steps of duration $\Delta t = 0.01$ (using a discrete version of the no-flux boundary conditions). Rather than using (7) to compute the concentration of ECs at lattice points, the factors P_0, P_1, P_2, P_3 and P_4 , whose formulae are given in (9)–(13), are used to determine the probability that an individual EC makes a move on the square lattice.

$$P_0 = 1 - \frac{4\Delta t D}{h^2} - \frac{\Delta t \chi c_{l,m}^t}{h^2} (c_{l+1,m}^t + c_{l-1,m}^t - 4c_{l,m}^t + c_{l,m+1}^t + c_{l,m-1}^t) - \frac{\Delta t \rho}{h^2} (f_{l+1,m}^t + f_{l-1,m}^t - 4f_{l,m}^t + f_{l,m+1}^t + f_{l,m-1}^t) \quad (9)$$

$$P_1 = \frac{\Delta t D}{h^2} - \frac{\Delta t}{4h^2} (\chi c_{l,m}^t (c_{l+1,m}^t - c_{l-1,m}^t) + \rho (f_{l+1,m}^t - f_{l-1,m}^t)) \quad (10)$$

$$P_2 = \frac{\Delta t D}{h^2} + \frac{\Delta t}{4h^2} (\chi c_{l,m}^t (c_{l+1,m}^t - c_{l-1,m}^t) + \rho (f_{l+1,m}^t - f_{l-1,m}^t)) \quad (11)$$

$$P_3 = \frac{\Delta t D}{h^2} - \frac{\Delta t}{4h^2} (\chi c_{l,m}^t (c_{l,m+1}^t - c_{l,m-1}^t) + \rho (f_{l,m+1}^t - f_{l,m-1}^t)) \quad (12)$$

$$P_4 = \frac{\Delta t D}{h^2} + \frac{\Delta t}{4h^2} (\chi c_{l,m}^t (c_{l,m+1}^t - c_{l,m-1}^t) + \rho (f_{l,m+1}^t - f_{l,m-1}^t)) \quad (13)$$

At each time step, each EC either remains at its location (l, m) , or moves to $(l-1, m)$, $(l+1, m)$, $(l, m-1)$ or $(l, m+1)$ according to the probabilities $\hat{P}_0, \hat{P}_1, \hat{P}_2, \hat{P}_3$ and \hat{P}_4 , where $\hat{P}_j = P_j / (P_0 + P_1 + P_2 + P_3 + P_4)$ for $j = 0, 1, 2, 3, 4$. To simulate this, a uniform random number $u \sim \mathcal{U}_{[0,1]}$ is drawn for each tip EC at each time-step. If $u \in [0, \hat{P}_0)$ then $s^{t+\Delta t} = s^t$, if $u \in [\hat{P}_0, \hat{P}_1)$, the EC moves left, and so on. If this procedure prescribes a move outside of the domain \mathcal{I} , that tip EC terminates and is not considered for any further moves.

A.2. Stokes-Lauffenberger (SL) model

In the SL model (Stokes et al., 1991), tip ECs can move in any direction (off-lattice). The movement of each tip EC is governed by (14), a two-dimensional Stochastic Differential Equation (SDE) for EC velocity $\mathbf{V}(t)$. At discrete time-steps, Equation (14) is solved using the Euler-Maruyama method and Equation (15) is integrated via the forward Euler method to give EC position.

$$d\mathbf{V}(t) = -\beta\mathbf{V}(t)dt + \sqrt{\sigma}d\mathbf{W}(t) + \Psi(t)dt \quad (14)$$

$$s(t) = \int_0^t \mathbf{V}(\tau)d\tau \quad (15)$$

In the SL model, the parameter β models the cell's resistance to movement, $\mathbf{W}(t)$ is a two-dimensional Wiener process (in which increments $\mathbf{W}(t+\Delta t) - \mathbf{W}(t)$ are independently and normally distributed), and the parameter σ represents an EC's tendency to deviate from its current direction. $\Psi(t) = \kappa \nabla c \sin \left| \frac{\psi}{2} \right|$ models movement due to chemotaxis, where c is the VEGF concentration and ψ is the angle that the current velocity $\mathbf{V}(t)$ makes with the direction of steepest increase in c . Given $c(x, y) = y$, the direction of increasing VEGF concentration is always $(0, 1)^T$, which simplifies the chemotaxis term. The parameter κ measures the strength with which the EC velocity re-orientates up spatial gradients of c (towards the tumour). ECs initialised along the bottom edge of the square domain are assigned a small initial velocity in the y -direction: $v^0 = (0, 0.5)^T$. The discretisation (16) of Equations (14) and (15) is then used to simulate EC velocity v^t and position s^t at each time-step.

$$v^{t+\Delta t} = (1 - \beta\Delta t)v^t + \sigma\Delta t\varepsilon + \kappa\Delta t\sqrt{\frac{1 - v_y^t/\|v^t\|}{2}}, \quad s^{t+\Delta t} = s^t + \Delta t v^{t+\Delta t} \quad (16)$$

The quantity $\varepsilon \sim \mathcal{N}(0, 1)$ is drawn from a two-dimensional normal distribution so that the random velocity vector has variance σ in each direction. Given the initial velocity, each tip EC in the SL model moves a distance 0.05 (in the vertical direction) in the first time-step, which is the same as one grid-space in the AC model. To ensure that ECs in the SL model move roughly this distance in every time-step (and the velocity does not grow exponentially), we fix $\beta = 0.8/\Delta t$.

A.3. Plank-Sleeman (PS) model

The PS model (Plank and Sleeman, 2004) assigns a constant speed \hat{s} to each EC and varies the angle ϕ that a tip EC's velocity vector makes with the horizontal direction. An EC's position $s^t = (s_x^t, s_y^t) \in \mathcal{I}$ is then modelled by the system of ordinary differential equations:

$$\frac{ds_x^t}{dt} = \hat{s} \cos \phi, \quad \frac{ds_y^t}{dt} = \hat{s} \sin \phi. \quad (17)$$

The movement angle is assumed to be independent of speed and position, and may be viewed as a random walk on the unit circle. At each time-step, a tip EC may turn clockwise or counter-clockwise through a fixed angle $\hat{\phi}$ or it may continue in the same direction. Given an initial movement angle ϕ_0 , the movement angle of each EC after n time steps is determined by transition probabilities $\hat{\tau}_n^\pm$. The transition probabilities are derived from the mean turning rate $\mu(\phi)$, which is given by:

$$\mu(\phi) = -d_c |\nabla c| \sin(\phi - \phi_c). \quad (18)$$

In Equation (18), the turning coefficient d_c determines how often an EC angle re-orientates its movement angle towards the direction of increasing VEGF concentration $\phi_c = \pi/2$. Let $\hat{\tau}_n^\pm$ denote the probability that an EC rotates through a fixed angle of $\pm\hat{\phi}$ on the n -th time-step. It can be shown (Plank and Sleeman, 2004) that, if the mean turning rate is defined by (18), then $\hat{\tau}_n^\pm$ are given by:

$$\hat{\tau}_n^\pm = 2\lambda \frac{\tau\left((n \pm \frac{1}{2})\hat{\phi}\right)}{\tau\left((n + \frac{1}{2})\hat{\phi}\right) + \tau\left((n - \frac{1}{2})\hat{\phi}\right)}, \text{ where } \tau(\phi) = \exp\left(\frac{d_c}{D_r} \cos(\phi - \phi_c)\right) \text{ and } \lambda = D_r/\hat{\phi}^2. \quad (19)$$

Choosing a random number $u \sim \mathcal{U}_{[0,1]}$ from the standard uniform distribution, an EC turns anticlockwise through an angle of $\hat{\phi}$ if $u \in [0, \hat{\tau}_n^+ \Delta t)$, clockwise through an angle of $\hat{\phi}$ if $u \in [\hat{\tau}_n^+ \Delta t, 2\lambda \Delta t)$, and continues in its current direction otherwise. Using this rule to generate movement angles ϕ_t at time steps $t \in \mathcal{T}$, the position $s^t = (s_x^t, s_y^t) \in [0, 1]^2$ of an individual EC is determined by solving the ODEs (17) using the forward Euler method:

$$s_x^{t+\Delta t} = s_x^t + \hat{s}\Delta t \cos \phi_t, \quad s_y^{t+\Delta t} = s_y^t + \hat{s}\Delta t \sin \phi_t. \quad (20)$$

Choosing a speed of $\hat{s} = 0.05$ ensures that, in the PS model, ECs move the same distance during each time-step as ECs in the AC model, and a similar distance as ECs in the SL model.

A.4. Anastomosis, branching and termination rules for all models

The previous sections describe how each model determines an EC's new position $s^{t+\Delta t}$ from its current position s^t . The rules for EC branching, anastomosis and termination are the same for all three models.

An active EC at location s^t may bifurcate into two separate EC. After branching, the original tip EC continues to move as instructed by its model, and a new variable s^t is created at the branch point to represent a new tip EC, which thereafter moves independently. We denote by a_{br} the minimum age that a EC must reach before it can be considered for a bifurcation. A second branching parameter, c_{br} , defines the minimum concentration of VEGF that must be present at s^t in order for the tip EC to bifurcate. In all models, EC bifurcate into two separate EC as soon as both the minimum age for branching and the VEGF threshold for branching have been exceeded. Since we initiate multiple ECs in the simulation domain, it is possible that a move may result in a tip EC crossing the path of an existing stalk EC (or colliding with another tip EC). In such cases, all three models assume an anastomosis event occurs. In the AC model, if a move requires an EC to move into a grid position that is already occupied, the EC does not make this move; it is terminated and not considered for any subsequent moves. In the SL and PS models, if a move requires an EC to cross an existing EC path, the EC terminates at the intersection of the proposed move and the existing path. If a model's movement rule specifies a new position $s^{t+\Delta t}$ which is outside the simulation domain \mathcal{I} , that tip EC terminates and is not considered for further movement.

Figure 4 gives an illustration of EC movement rules, including examples of anastomosis, which can lead to loops in the simulated network.

A.5. Model simulation, schematic and parameters

We fix all but four parameter values in each angiogenesis model and simulate the movement of tip ECs for 400 time-steps. See Table 1 for a list of the parameters that vary and their ranges. See also the schematic at the top of Figure 4 for an illustration of model movement rules.

We convert data simulated from the three angiogenesis models into a common format for data analysis. The common format should retain the spatial structure of simulated networks while discretising them to allow the computation of spatially-averaged and topological statistics. Considering a simulated network at its final timestep, we overlay a regular square grid of $K = 200$ points spaced in intervals of $h = 0.05$ in co-ordinate directions over the simulation domain. We then say an *angiogenesis dataset* is the point-cloud consisting of the (x, y) locations of ECs within this discretised image. The AC model simulates data on such a grid already, and we convert data from the SL and PS models into the common format by populating those grid squares which intersect the (off-lattice) paths traced by tip ECs. This common format will also be applied to observed data and ensures that summary statistics give a fair comparison between data simulated by each model.




Fig. 4: Schematic showing tip EC movement rules in the Anderson-Chaplain (AC), Stokes-Lauffenberger (SL) and Plank-Sleeman (PS) models. In the AC model, tip ECs move on a grid according to probabilities $\hat{P}_j = P_j / (P_0 + P_1 + P_2 + P_3 + P_4)$ for $j = 0, 1, 2, 3, 4$. Higher values of the chemotaxis and haptotaxis parameters induce a bias into those probabilities which specify a movement towards increasing concentrations of VEGF and fibronectin respectively. The SL model updates the velocity v_i^t of the EC at location s^t as a weighted sum of the current velocity w , randomness r and chemotaxis c . Randomness and chemotaxis parameters moderate the weight of the corresponding terms when updating the velocity. The PS model rotates the movement angle ϕ between an EC's velocity vector and the horizontal by $\hat{\phi}$ with transition probabilities $\hat{\tau}_n^+$ and $\hat{\tau}_n^-$ and ECs move a fixed distance \hat{s} in the new direction at each time-step. A turning rate parameter moderates how often the EC's angle of movement updates, and a turning bias parameter changes how likely such an update is to favor the direction of increasing VEGF concentration. All models use the same rules for branching—a tip EC at location s^t bifurcates into two tip ECs that move independently when t is greater than the minimum age for branching parameter a_{br} and the concentration of VEGF at s^t is greater than the VEGF threshold for branching parameter c_{br} . If a movement rule would cause an EC to move into a location already occupied by an EC, that EC instead anastomoses and is considered for no further movement. Created in

<https://BioRender.com>

Anderson-Chaplain (AC) model			Stokes-Lauﬀenberger (SL) Model			Plank-Sleeman (PS) Model		
parameters			parameters			parameters		
Name and Symbol	Range	Notes	Name and Symbol	Range	Notes	Name and Symbol	Range	Notes
Chemotaxis χ	[0, 0.5]	Taken from (Nardini <i>et al.</i> , 2021)	Chemotaxis κ	[0, 100]	The chemotaxis term in (16) has the same order of magnitude as the current velocity	Turning bias d_c	[30, 150]	30 gives no turning bias, 150 almost always favours turning towards the VEGF source
Haptotaxis ρ	[0, 0.5]	Taken from (Nardini <i>et al.</i> , 2021)	Randomness σ	[0, 100]	The randomness term in (16) has the same order of magnitude as the current velocity	Turning rate D_r	[0, 30]	A turning rate of 30 means EC change direction every other time-step
Minimum age for branching a_{br}	[0, 1]	ECs can bifurcate after between 0 and 100 time-steps.	Minimum age for branching a_{br}	[0, 1]	ECs can bifurcate after between 0 and 100 time-steps.	Minimum age for branching a_{br}	[0, 1]	ECs can bifurcate after between 0 and 100 time-steps.
VEGF threshold for branching c_{br}	[0, 1]	The concentration of VEGF varies in \mathcal{I} from 0 when $y = 0$ to 1 when $y = 1$.	VEGF threshold for branching c_{br}	[0, 1]	The concentration of VEGF varies in \mathcal{I} from 0 when $y = 0$ to 1 when $y = 1$.	VEGF threshold for branching c_{br}	[0, 1]	The concentration of VEGF varies in \mathcal{I} from 0 when $y = 0$ to 1 when $y = 1$.

Table 1. The parameters inferred in each model and their values/ranges.

B. Computation of Summary Statistics

B.1. Spatially-Averaged Statistics

An angiogenesis dataset \mathcal{D} is a point-cloud consisting of (x, y) grid locations which contain simulated EC at the final simulation timestep. We compute the mean, standard deviation, minimum, maximum, range, interquartile range, and the 10th, 25th, 75th and 90th percentiles of the x and y -coordinates of points in \mathcal{D} . Concatenating these values gives a list of 20 spatially-averaged statistics, 10 in the horizontal (x) coordinate direction and 10 in the vertical (y) coordinate direction.

B.2. Extended Persistent Homology

We detect and quantify topological structure in angiogenesis datasets using Extended Persistent Homology (EPH). Here, we give details of how we compute and vectorize EPH, giving a brief overview of how it arises from Persistent Homology (PH) and outlining our reasons for using EPH instead of PH.

To compute PH, spatial data is first converted to a nested sequence of simplicial complexes $\{\Sigma_k\}_{k=0}^K$ known as a filtration. A simplicial complex is a collection of vertices (0-simplices), where subsets of vertices can be connected by edges (1-simplices), triangles (2-simplices), tetrahedra (3-simplices) and their higher dimensional analogues. A face of a simplex is defined as a subset of its vertices along with the simplices that are connected to them. Simplicial complexes satisfy the property that a face of any simplex, or the intersection of two simplices, is also a simplex in the complex.

The dimension- p PH of the filtration is then the sequence (21) of \mathbb{F} -vector spaces $H_p(\Sigma_k)$ (we fix $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$) together with maps induced by the inclusion maps between the simplicial complexes. The basis elements of $H_p(\Sigma_k)$ correspond to p -dimensional holes in Σ_k (Carlsson, 2009), also referred to as topological features of the simplicial complex. Since each $\Sigma_k \subset \Sigma_{k+1}$, inclusion maps $\iota_k : \Sigma_k \rightarrow \Sigma_{k+1}$ map simplices in Σ_k to their counterparts in the larger Σ_{k+1} . The inclusion maps ι_k then induce linear maps $\iota_k^* : H_p(\Sigma_k) \rightarrow H_p(\Sigma_{k+1})$ in sequence (21) (by functoriality) which allow topological features to be tracked through the filtration.

$$H_p(\Sigma_0) \xrightarrow{\iota_0^*} H_p(\Sigma_1) \xrightarrow{\iota_1^*} \dots \xrightarrow{\iota_{k-1}^*} H_p(\Sigma_k) \xrightarrow{\iota_k^*} \dots \xrightarrow{\iota_{K-2}^*} H_p(\Sigma_{K-1}) \xrightarrow{\iota_{K-1}^*} H_p(\Sigma_K) \quad (21)$$

The Structure Theorem of PH (Zomorodian and Carlsson, 2005) states that sequence (21) uniquely decomposes into a direct sum of interval modules $I_{b,d}$, which are sequences of \mathbb{F} -vector spaces (22), where η_k^* are identity maps when $b \leq k < d$ and zero maps otherwise.

$$0 \xrightarrow{\eta_0^*} 0 \dots 0 \xrightarrow{\eta_{b-1}^*} \mathbb{F} \xrightarrow{\eta_b^*} \mathbb{F} \dots \mathbb{F} \xrightarrow{\eta_{d-1}^*} 0 \dots 0 \xrightarrow{\eta_{K-1}^*} 0 \quad (22)$$

Persistence pairs (b, d) that define interval modules (22) correspond to topological features in the filtration. The birth b is the first index k for which the corresponding p -dimensional hole appears in the filtration (and the index k in (22) for which the corresponding topological feature is not in the image of η_{k-1}^*). The death d is the filtration index k where the p -dimensional hole is filled in by additional simplices (and the index k in (22) for which the corresponding topological feature is mapped to 0 by η_{k-1}^*). The persistence $d - b$ measures how long the corresponding topological feature persists in the filtration. Persistence pairs may be plotted as points in birth-death coordinates in a Persistence Diagram (PD). PDs are stable (Chazal et al., 2009) to small perturbations in the underlying point-cloud, making them useful topological summaries of spatial data.

Nardini et al. (2021) applied PH to study angiogenesis datasets \mathcal{D} by constructing a sweeping-plane filtration as follows. First, \mathcal{D} is converted to a simplicial complex Σ . Each point (x, y) in the angiogenesis dataset is represented in Σ by a vertex (0-simplex). If two vertices represent two grid locations in \mathcal{D} which are adjacent (in the Moore neighbourhood), they are connected with edges (1-simplices). Collections of three edges are connected with triangle (2-simplex) if each pair of edges shares a vertex. A function $f : \Sigma \rightarrow \mathbb{R}$ is then defined such if a vertex u represents a point $(x, y) \in \mathcal{D}$, $f(u) = y$. The value of f on other simplices within Σ is then simply the maximum of the values of f on their vertices. The sweeping-plane filtration is then defined as the sublevel sets $\Sigma_k = f^{-1}(-\infty, k/K]$, where $k = 0, 1, \dots, K$ and K is the resolution of the image (e.g. $K = 200$).

Using this filtration, persistence pairs quantify the location of components and loops (measured in the vertical (y) coordinate direction). However, some persistence pairs computed using this method have infinite persistence, since topological features often persist throughout the entirety of the sequence (21) (and $d = \infty$). For example, all loops persistent infinitely in this filtration. In these cases, the size of the p -dimensional holes is not recovered by the persistence pair (b, d) . Here, we wish to quantify both the size and location of topological features in angiogenesis datasets, and we wish to precisely measure branch points, loops, anastomoses and components, which PH with the above sweeping-plane filtration is not able to do. We therefore turn to Extended Persistent Homology (EPH).

To compute EPH, one defines Σ^k as the simplicial complex containing those simplices in Σ_K but not Σ_k . With the sweeping-plane filtration above, Σ^k are superlevel sets $f^{-1}[k/K, \infty)$. The relative homology $H_p(\Sigma_K, \Sigma^k)$ then quantifies topological features in the quotient complex Σ_K/Σ^k . Since $\Sigma^{k+1} \supset \Sigma^k$, a quotient map $q_k : \Sigma^{k+1} \rightarrow \Sigma^k$ may be defined for $k = K-1, \dots, 0$. q_k maps all simplices in Σ^{k+1} but not Σ^k to a single point and is the identity map when restricted to Σ^k . The quotient maps q_k induce linear maps in sequence (23) which track topological features through the relative homology groups.

$$H_p(\Sigma_K, \Sigma^K) \xrightarrow{q_{K-1}^*} H_p(\Sigma_K, \Sigma^{K-1}) \xrightarrow{q_{K-2}^*} \dots \xrightarrow{q_k^*} H_p(\Sigma_K, \Sigma^k) \xrightarrow{q_{k-1}^*} \dots \xrightarrow{q_1^*} H_p(\Sigma_K, \Sigma^1) \xrightarrow{q_0^*} H_p(\Sigma_K, \Sigma^0) \quad (23)$$

Since $H_p(\Sigma_K) = H_p(\Sigma_K, \Sigma^K)$, the sequences (21) and (23) may be concatenated to give a single sequence (of length $2k + 1$), which is known as the dimension- p EPH of the filtration. The structure theorem applies to this sequence, meaning that topological features may

be extracted as persistence pairs (b, d) . Intuitively, the sweeping-plane filtration scans through Σ from bottom to top and (21) detects topological features which are found below k/K . The complexes Σ_K/Σ^k then scan back down from top to bottom, collapsing all simplices above k/K to a single point, with the sequence (23) detecting topological features in the resulting complexes. In particular, all simplices merge in Σ_K/Σ^0 , meaning that no persistence pairs extracted from (21)–(23) have $d = \infty$.

Persistence pairs (b, d) representing topological features which are born and die within the sequence (21) are called ordinary persistence pairs. The sweeping-plane filtration detects such topological features in dimension $p = 0$ when a path of ECs anastomoses with another (without forming a loop). Topological features which are born and die in the sequence (23) are called relative persistence pairs. In dimension $p = 1$ such topological features result from branch points where an EC bifurcates in two separate paths which do not later anastomose to form loops. Those topological features which are born in (21) but die in (23) are extracted as extended persistence pairs. If $b < d$ and $p = 0$, extended persistence pairs represent connected components—distinct EC paths which remain separated. If $b > d$ and $p = 1$, extended persistence pairs represent loops in the simulated vasculature.

We use two functions f to create two sweeping-plane filtrations from an angiogenesis dataset \mathcal{D} . Creating a simplicial complex Σ as above, we define $f_{\text{horizontal}}, f_{\text{vertical}} : \Sigma \rightarrow \mathbb{R}$ on vertices u representing points $(x, y) \in \mathcal{D}$ by $f_{\text{horizontal}}(u) = x$, $f_{\text{vertical}}(u) = y$, and on other simplices as before. Persistence pairs resulting from the sequences (21)–(23) then represent the size and location of topological features measured in the horizontal (x) and vertical (y) coordinate directions. Persistence pairs may be plotted on an Extended Persistence Diagram (EPD), an example of which is given in Figure 5.

B.3. Topological Statistics

We use Persistence Images and Persistence Statistics to transform EPDs into fixed-length vectors amenable to further analysis.

Persistence images are stable vector representations of persistence diagrams (Adams *et al.*, 2017), and they were used by Thorne *et al.* (2022) as summary statistics to infer two parameters in the AC model. To obtain persistence images, persistence pairs (b, d) are first transformed to the modified pair $(b, d - b)$. A persistence surface is then the function $\Xi : \mathbb{R}^2 \rightarrow \mathbb{R}$ (24), where the sum is taken over all modified pairs.

$$\Xi(x, y) = \sum_{(b, d-b)} w(b, d-b) \Phi_{(b, d-b)}(x, y). \quad (24)$$

Here, w is a weighting function on the modified persistence pairs, and we use the standard choice $w(b, d-b) = d-b/\max(d-b)$, which divides the persistence by the maximum persistence across all persistence pairs. The function $\Phi_{(b, d-b)}$ must be a differentiable probability distribution with mean $(b, d-b)$, and we take the standard Gaussian distribution with variance 0.1. A persistence image is then a collection of integrals (25) of Ξ over discretised regions $\mathcal{R} \subset \mathbb{R}^2$:

$$\Gamma(\Xi)_{\mathcal{R}} = \int \int_{\mathcal{R}} \Xi dy dx. \quad (25)$$

The regions \mathcal{R} must be chosen so that the integrals reflect the distribution of persistence pairs used to construct Ξ . The persistence pairs we compute from angiogenesis datasets with sweeping-plane filtrations have birth and death values between $k = 0$ and $k = 200$. We therefore take regions $R_{i,j} = [20i, 20(i+1)] \times [20j, 20(j+1)]$ for $i, j = 0, 1, \dots, 9$ to generate a total of 100 integrals which constitute a persistence image. For each angiogenesis dataset we compute eight such persistence images—one for each of the four types of extended persistence pair described in Section B.2 computed using each sweeping-plane filtration (horizontal and vertical), giving a total of 800 topological statistics.

Persistence statistics are simple summaries of EPDs that have been found to perform well on well-known classification tasks (Ali *et al.*, 2023). To obtain persistence statistics, we consider the births b , deaths d , persistences $d-b$ and midpoints $(b+d)/2$ of persistence pairs (b, d) in an EPD, and compute the mean, standard deviation, median, interquartile range, full range, and the 10th, 25th, 75th and 90th percentiles of each of these descriptors. We therefore compute a total of 36 persistence statistics for each of the four types of persistence pairs described in Section B.2 for each sweeping-plane filtration (horizontal and vertical), for a total of 288.

B.4. Example Computation

We compute spatially-averaged and topological statistics from angiogenesis datasets and concatenate these into a long-list of summary statistics. Each angiogenesis dataset yields 20 spatially-averaged statistics and $800 + 288$ topological statistics, which we combine into a vector of length 1108. An example of the computation of spatially-averaged statistics and EPDs for a simple angiogenesis dataset is given in Figure 5.

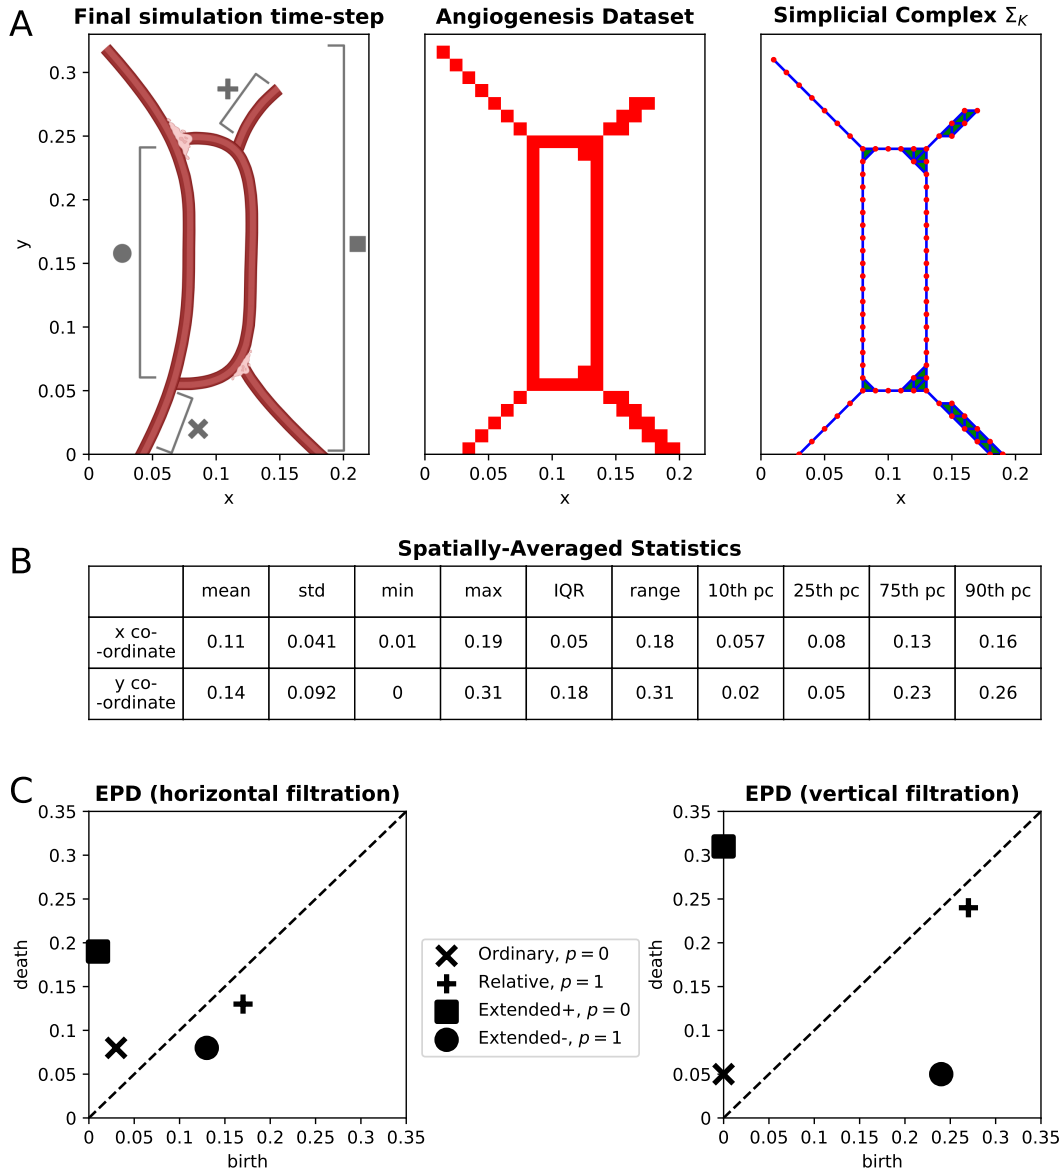


Fig. 5: An example model simulation and computation of spatially-averaged statistics and EPDs. A) Data pre-processing. We display the final time-step of a simple simulation in which two ECs migrate from the bottom of the simulation domain, branch, anastomose and form a loop in their path upward towards the tumour. We convert the final simulation time-step into an angiogenesis dataset by overlaying a grid of 200×200 pixels onto the simulation domain and colouring pixels which correspond to the location of simulated ECs. Simplicial complexes Σ_k contain a vertex for each pixel location containing an EC whose x co-ordinate (or y coordinate) is less than or equal to $k/200$ —we plot Σ_K where $K = 200$. B) Spatially-averaged statistics. Spatially-averaged statistics consider the (x, y) co-ordinates of pixel locations which contain ECs. We compute the mean, standard deviation, minimum, maximum, interquartile range, range, 10th percentile, 25th percentile, 75th percentile and 90th percentile in each co-ordinate direction. C) EPDs. We compute EPDs to quantify the size and location of the loops (●), the branches (× and +), and components (■) at the final time-step, where the size and location are measured in either the horizontal (x) or vertical (y) direction. Computing EPH via the sequences (21) and (23) using the horizontal and vertical filtration gives two EPDs containing persistence pairs (b, d) which correspond to the birth and death of topological features within the concatenated sequence. The values b and d correspond to the location (in co-ordinate directions) of each feature and the persistence $d - b$ gives the size. As the final step in computing topological statistics, we vectorise EPDs into persistence images and persistence statistics according to Section B.3. The left panel of A was created in <https://BioRender.com>

C. ABC and Random Forest Methodology

We use Random Forests (Breiman, 2001) to choose informative summary statistics, which we then use to infer spatial parameters and select between spatial models using ABC. In this section, we give further details on each step described in Section 3 of the main text.

C.1. Step 1: Identify Informative Summary Statistics For Each Model

ABC requires a distance function $\nu : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ to measure the discrepancy between observed data and simulated data. In Section B we showed how to compute spatially-averaged and topological statistics from angiogenesis datasets. Combining all spatially-averaged and topological statistics into a single vector gives a total of $n_f = 1108$ features which may be computed from an angiogenesis dataset. We now identify an informative subset of these summary statistics which we will use to construct the distance function ν .

We train one regression Random Forest per parameter per angiogenesis model to learn the relationship between parameter values and summary statistics. By training one Random Forest per parameter, we aim to identify the summary statistics that capture the effect of varying each individual parameter's value on simulated data. To generate training data, we draw $n = 10,000$ parameter values θ_i uniformly from the parameter ranges given in Table 1. We generate an angiogenesis dataset corresponding to each parameter value by simulating each model up to time $t_{\text{final}} = 4$ and consider the angiogenesis dataset at its final time-step as described in Section A. We compute spatially-averaged and topological statistics to form a vector X_i of summary statistics corresponding to the parameter value θ_i , which together form training data \mathcal{X} and \mathcal{Y} .

To construct decision trees that make up each Random Forest, we take bootstraps of $n_{\text{samples}} = 2,000$ pairs (X_i, θ_i) sampled independently (with replacement) from the training data. Each decision tree is made up of internal nodes which repeatedly partition the training data and leaf nodes which contain $n_{\text{min}} = 5$ or fewer pairs from the bootstrap. At an internal node \mathcal{N} , $n_f/3$ features are randomly selected and considered to create a splitting rule. Splitting rules are conditions of the form $X_i^j > r$ for some co-variate j and splitting bound r . Pairs $(X_i, \theta_i) \in \mathcal{N}$ for which a condition is satisfied are allocated to the right daughter node \mathcal{N}_R and the others are passed to left daughter node \mathcal{N}_L . The bootstrap of the training data is repeatedly partitioned by internal nodes in this way until $n_{\text{min}} = 5$ or fewer pairs (X_i, θ_i) are allocated to a node, wherein it becomes a leaf node.

To decide the co-variate index j and splitting bound r used at each internal node, Random Forests consider a loss function with the general form of Equation (26). Random Forests choose co-variables and splitting bounds which minimise Δ_{loss} within each internal node.

$$\Delta_{\text{loss}} = \frac{|\mathcal{N}_L|}{|\mathcal{N}|} Q(\mathcal{N}_L) + \frac{|\mathcal{N}_R|}{|\mathcal{N}|} Q(\mathcal{N}_R) \quad (26)$$

$Q(\mathcal{N})$ is a measure of the impurity of samples allocated to node \mathcal{N} and $|\mathcal{N}|$ is the number of samples allocated to the node. As is common in regression Random Forests, we use the L2 impurity given by Equation (27) to decide splitting rules. The L2 purity measures the variance in parameter values from the mean $\bar{\theta}_{\mathcal{N}}$ among pairs allocated to the same daughter node.

$$Q(\mathcal{N}) = \sum_{\theta_i : (X_i, \theta_i) \in \mathcal{N}} (\theta_i - \bar{\theta}_{\mathcal{N}})^2 \quad (27)$$

To choose the number n_{tree} of decision trees to use in each Random Forest, we follow the advice in Raynal *et al.* (2018) and compute the out-of-bag mean square error of Random Forests constructed with different numbers of trees. We note in Figure 6 that the out-of-bag error decreases when adding additional trees, but that this improvement is small when using more than 100 trees in the Random Forest for each parameter in each model. We therefore use $n_{\text{tree}} = 100$.

To select a subset of the spatially-averaged and topological statistics for use in subsequent ABC algorithms, we consider the importance of each feature in the Random Forest for each parameter. The importance of a co-variate j in a trained Random Forest is defined as the mean decrease in impurity $Q(\mathcal{N}) - (Q(\mathcal{N}_L) + Q(\mathcal{N}_R))$ achieved by all internal nodes \mathcal{N} which use j in their splitting rule. Features with high importance are therefore those which are most effective in partitioning the training data, meaning they capture the effect of model parameters on model simulations.

Ranking the $n_f = 1108$ summary statistics by importance, we observe in Figure 7 that a small number of features carry most of the predictive power for each parameter in each model. To construct the distance function ν used in ABC, we wish to use those statistics which distinguish data simulated using different model parameters and omit those which do not. Therefore, we cycle through each parameter in each model and choose the top 25 most important summary statistics for each parameter that have not already been selected, giving a list of 100 features in total for each model. In general, we recommend computing the importance of a long-list of summary statistics and choosing n_s so that statistics with low importance are omitted.

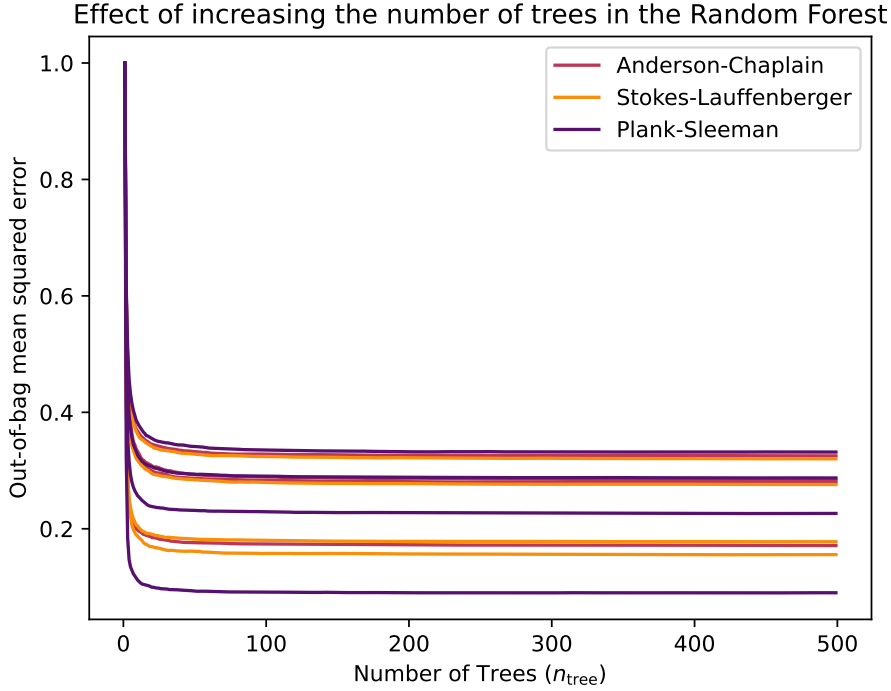


Fig. 6: The out-of-bag mean square error of regression Random Forests learning the relationship between summary statistics X_i and parameter values θ_i (scaled by its initial value). The out-of-bag mean square error initially decreases, but negligibly after $n_{\text{tree}} = 100$.

C.2. Step 2: Fit Each Model To The Observed Data

Using observed data \mathcal{D}^* , we approximate the parameter posterior $p(\Theta|\mathcal{D}^*)$ for each angiogenesis model. To construct the distance function ν required by ABC algorithms to measure the discrepancy between simulated and observed data, we follow Section C.1 and choose a subset of 100 summary statistics for each model from among the spatially-averaged and topological statistics. It is possible that the important summary statistics may be poorly scaled. For example, connected components are generally larger than loops in the angiogenesis datasets we consider. Small differences in the size and location of loops in angiogenesis datasets may be stronger predictors of parameter values than the size and locations of connected components, yet this may not be reflected in the distance function ν if the corresponding topological statistics are left unscaled. We therefore construct a scaling function which divides each summary statistic by the maximum absolute value of that summary statistic among the training data for each model. The ABC distance function is then $\nu(\mathcal{D}^*, \mathcal{D}_i) = \|x^* - x_i\|_2$ where \mathcal{D}^* is observed data, \mathcal{D}_i is simulated data, and x^* and x_i are the 100 summary statistics with the scaling applied. When \mathcal{D}^* comprises several instances of observed data, we average the distance ν using summary statistics computed from each.

Using this distance function, we use the ABC-SMC algorithm of Del Moral et al. (2012) to approximate the parameter posterior $p(\Theta|\mathcal{D}^*)$. ABC-SMC outputs a series of n_{pop} intermediate distributions, which correspond to a decreasing sequence of tolerances $\epsilon_0 > \dots > \epsilon_{n_{\text{pop}}} > 0$. Beginning with a population of N_{pop} parameters sampled from a prior distribution, the initial tolerance ϵ_0 is chosen so that a predetermined fraction $\alpha \in (0, 1)$ of parameters are accepted. The Effective Sample Size (ESS) of each population is a measure of the independence of parameters within a current population. Del Moral et al. (2012) computes the ESS of each subsequent population and chooses further tolerances such the ESS of each population decreases by α . We use $\alpha = 0.8$ and generate $n_{\text{pop}} = 50$ populations, each containing $N_{\text{pop}} = 1,000$ parameters, and take the final population as the approximate ABC-SMC posterior $p(\Theta|\mathcal{D}^*)$.

C.3. Step 3: Approximate The Model Posterior

We approximate the model posterior $p(m|\mathcal{D}^*)$ for observed data \mathcal{D}^* by training two more Random Forests, following the method of Pudlo et al. (2015). Using the training data (X_i, m_i) obtained from the training data used in Step 1 by replacing the parameter θ_i with the model index $m_i \in \{\text{AC}, \text{SL}, \text{PS}\}$, we train a classification Random Forest to learn the relationship between summary statistics X_i and model index m_i .

Step 1 chooses a subset of the spatially-averaged and topological statistics which are important for each for each model. Since the importance of each summary statistic to a model is determined by that model's training data, it is possible that different summary statistics are important for different models. Using summary statistics which are only informative for a single model to approximate the

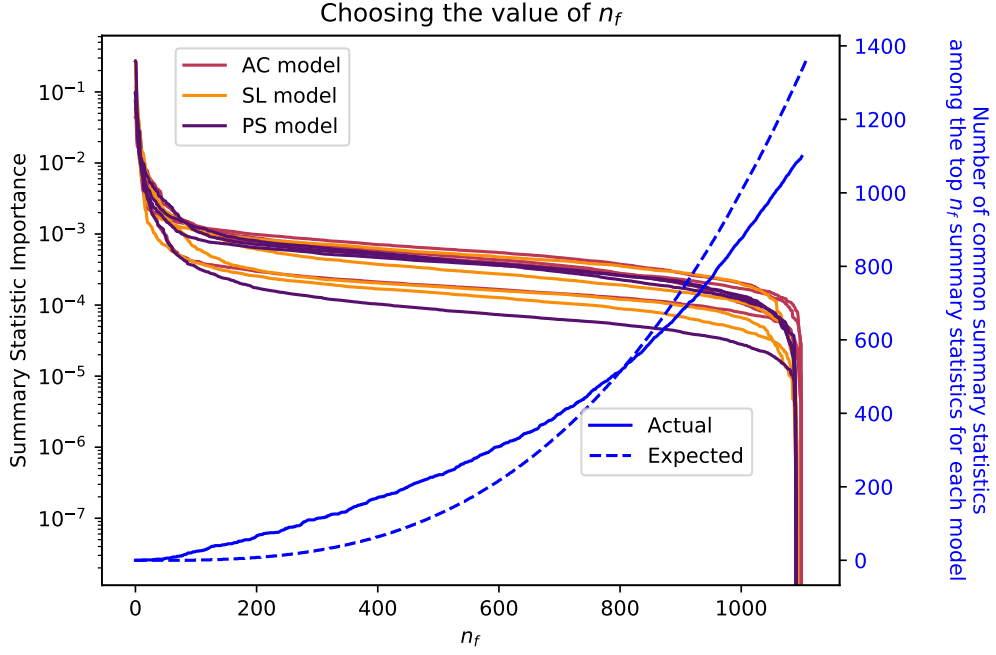


Fig. 7: Left axis: for each parameter in each model, we rank summary statistics by their importance and observe how importance decreases. Right axis: we compare how many summary statistics appear among the top n_s most important features for all models as n_s is increased and plot the expected number of common summary statistics if they were chosen randomly. Choosing $n_s = 100$ includes approximately the 25 most important summary statistics for each parameter. Selecting more summary statistics than this would include those which are of little use in capturing the effect of parameter values on simulated data. Choosing $n_s = 100$ also ensures that $\tilde{n}_s = 30$ summary statistics appear in the top n_s most important summary statistics for all models. If 100 summary statistics were chosen uniformly randomly for each model, the expected number that would be common to all three models is approximately 1.

model posterior may bias predictions in favour of that model. We therefore learn the relationship between summary statistics and model index using only those summary statistics which are important for all angiogenesis models.

Specifically, we modify X_i to contain only those $\tilde{n}_s \leq 100$ features which appear in the top 100 summary statistics for all three models. We found that $\tilde{n}_s = 30$ summary statistics appear in the top 100 for all angiogenesis models, and hence the X_i are modified to include only these entries. Figure 7 shows how the number of summary statistics that appear in the top n_s summary statistics for all models grows as n_s increases. We also plot the expected number of common summary statistics among the three models if summary statistics were selected randomly. n_s should be large enough to include some common summary statistics that important for all models, but not so large as to allow the possibility that some summary statistics are chosen for all three models by random chance. The expected number of common summary statistics if $n_s = 100$ are chosen randomly for each model is approximately 1, which suggests that the $\tilde{n}_s = 30$ common summary statistics are indeed informative for all three models, and were not chosen by random chance.

As in Step 1, we bootstrap $n_{\text{samples}} = 2,000$ pairs (X_i, m_i) to train $n_{\text{tree}} = 100$ decision trees, randomly choosing $\tilde{n}_s/3$ summary statistics at each internal node to consider for splitting rules, and terminating at leaf nodes only when $n_{\text{min}} = 5$ samples remain. To decide on the co-variate index j and splitting bound r in the partition rule $X_i^k < r$ at each internal node, we use the Gini impurity defined in Equation (28). For a node \mathcal{N} , let p_i be the proportion of pairs (X_i, m_i) allocated to \mathcal{N} . The Gini impurity is a measure of how many different model indices are allocated to node \mathcal{N} , which the splitting rule minimises.

$$Q(\mathcal{N}) = \sum_{m_i \in \{\text{AC}, \text{SL}, \text{PS}\}} p_i(1 - p_i) \quad (28)$$

The trained classification Random Forest provides a prediction $RF(X^*)$ of the model m^* which generated observed data \mathcal{D}^* (which we aggregate when \mathcal{D}^* contains multiple instances of observed data). The Random Forest also provides an out-of-bag error for each pair (X_i, m_i) , which is the proportion of those decision trees which did not use this pair in their training bootstrap which predict the incorrect model index m_i . Following Pudlo *et al.* (2015), we train another regression Random Forest using the same training data to learn the relationship between summary statistics X_i and this mis-classification error rate $p(RF(X_i) \neq m_i)$ of the first Random Forest. Predicting the out-of-bag error $p(RF(X^*) \neq m^*)$ for the observed data \mathcal{D}^* gives an estimate $1 - p(RF(X^*) \neq m^*)$ for the model posterior $p(m = m^* | \mathcal{D}^*)$.

D. Toy Model Example

D.1. Two Two-Parameter Toy Models

We verify our pipeline for parameter inference and model selection on two toy models which can simulate angiogenesis datasets similar to the simple vascular network in Figure 5. While the likelihood functions for the angiogenesis models analysed in the main text are unavailable, we construct simple toy models from which we can derive a likelihood function. Using appropriate priors and Bayes' rule, we can then derive the exact parameter posterior for each toy model, and compare it to the approximate parameter posterior produced using our pipeline (Steps 1-2 of Section 3 of the main text). We then test our method of model selection (Step 3 of Section 3 of the main text) on the two toy models.

Both toy models initialise two tip EC along the bottom of a the square domain $\mathcal{I} = [0, 1]^2$, which move diagonally upwards. The tip ECs anastomose, branch, and then anastomose again, forming a central loop before reaching the VEGF source at the top of the domain. We derive two stochastic spatial models from this simple construction by parametrising the size and shape of the central loop. The first model, Toy-Circle (TC), has two parameters r and c which define the centre and radius of a circular loop in the middle of the domain. Given a parameter pair (r, c) , the TC model produces a blood vessel network whose central loop is a circle of radius $r + \varepsilon$ and centre $(0.5, c + \varepsilon)$, where each ε is identically and independently drawn from $\mathcal{U}_{[-0.03, 0.03]}$, the uniform distribution between -0.03 and 0.03 . The randomly sampled ε are added to each parameter to introduce random variation into the model which is simple enough to allow derivation of the model's likelihood function. The second model, Toy Ellipse (TE), uses a similar construction to produce a vascular network. However, the TE model uses parameters a and b , to define the size of the horizontal radius and vertical radius of an ellipse which makes up the loop in the centre of the domain. Given a parameter pair (a, b) , the TE model produces a blood vessel network whose central loop is an ellipse centred at $(0.5, 0.5)$ with horizontal radius $a + \varepsilon$ and vertical radius $b + \varepsilon$. The noise ε is sampled in the same way as in the TC model. Each toy model, including the range of each parameter, is summarised in Table 2.

D.2. Computing The Exact Likelihood and Posterior

Suppose an angiogenesis dataset \mathcal{D}_i is simulated from the TC model using parameters (r, c) and has radius r_i and centre c_i . The likelihood of \mathcal{D}_i is the product of uniform distributions: $p(\mathcal{D}_i | r, c) = p(r_i, c_i | r, c) = \mathcal{U}_{[r-0.03, r+0.03]}(r_i) \times \mathcal{U}_{[c-0.03, c+0.03]}(c_i)$. Assume that uniform priors $p(r) = \mathcal{U}_{[r_{\min}, r_{\max}]}(r)$ and $p(c) = \mathcal{U}_{[c_{\min}, c_{\max}]}(c)$ are used for r and c respectively, with the maximum and minimum value of each parameter taken from Table 2. Given observed data $\mathcal{D}^* = \{\mathcal{D}_1^*, \dots, \mathcal{D}_{n^*}^*\}$ consisting of n^* angiogenesis datasets whose central loops have centres $c_1^*, \dots, c_{n^*}^*$ and radii $r_1^*, \dots, r_{n^*}^*$, the parameter posterior can be computed exactly by (29)–(33).

$$p(r, c | \mathcal{D}^*) = \prod_{i=1}^{n^*} p(r, c | \mathcal{D}_i^*) \propto \prod_{i=1}^{n^*} p(\mathcal{D}_i^* | r, c) \times p(r, c) \quad (29)$$

$$= \prod_{i=1}^{n^*} p(r_i^*, c_i^* | r, c) \times p(r) \times p(c) \quad (30)$$

$$= \prod_{i=1}^{n^*} [\mathcal{U}_{[r-0.03, r+0.03]}(r_i^*) \times \mathcal{U}_{[c-0.03, c+0.03]}(c_i^*)] \times \mathcal{U}_{[r_{\min}, r_{\max}]}(r) \times \mathcal{U}_{[c_{\min}, c_{\max}]}(c) \quad (31)$$

$$= \prod_{i=1}^{n^*} [\mathcal{U}_{[r_i^*-0.03, r_i^*+0.03]}(r) \times \mathcal{U}_{[c_i^*-0.03, c_i^*+0.03]}(c)] \times \mathcal{U}_{[r_{\min}, r_{\max}]}(r) \times \mathcal{U}_{[c_{\min}, c_{\max}]}(c) \quad (32)$$

$$= \mathcal{U}_{[\max\{r_{\min}, \max_{i=1}^{n^*} r_i^* - 0.03\}, \min\{r_{\max}, \min_{i=1}^{n^*} r_i^* + 0.03\}]}(r) \times \mathcal{U}_{[\max\{c_{\min}, \max_{i=1}^{n^*} c_i^* - 0.03\}, \min\{c_{\max}, \min_{i=1}^{n^*} c_i^* + 0.03\}]}(c) \quad (33)$$

Line (29) uses Bayes' rule and factors out the evidence $p(\mathcal{D}_i^*)$. We assume throughout that r_i^* and c_i^* have values within the ranges given in Table 2 (plus or minus 0.03) and are within 0.06 of each other (otherwise the likelihood and posterior are both 0). For the TE model, assuming that observed data $\mathcal{D}^* = \{\mathcal{D}_1^*, \dots, \mathcal{D}_{n^*}^*\}$ is a collection of angiogenesis datasets where the central loops are ellipses with horizontal radii $a_1^*, \dots, a_{n^*}^*$ and vertical radii $b_1^*, \dots, b_{n^*}^*$, a similar computation gives the parameter posterior $p(a, b | \mathcal{D}^*)$ as (34).

$$p(a, b | \mathcal{D}^*) = \mathcal{U}_{[\max\{a_{\min}, \max_{i=1}^{n^*} a_i^* - 0.03\}, \max\{a_{\max}, \min_{i=1}^{n^*} a_i^* + 0.03\}]}(a) \times \mathcal{U}_{[\max\{b_{\min}, \max_{i=1}^{n^*} b_i^* - 0.03\}, \min\{b_{\max}, \min_{i=1}^{n^*} b_i^* + 0.03\}]}(b) \quad (34)$$

D.3. Parameter Inference and Model Selection

We generate four synthetic test-cases for each toy model by choosing parameter pairs which cover a range of circles/ellipses. We generate $n^* = 2$ simulations at each test parameter pair and use them as observed data \mathcal{D}^* . Following Steps 1-2 of Section 3 of the main text, we approximate the parameter posteriors $p(r, c | \mathcal{D}^*)$ for the TC model and $p(a, b | \mathcal{D}^*)$ for the TE model in each test-case.

In Figure 8, we plot the true posterior, calculated using Equations (33) and (34), as a light blue square, and the approximate posterior resulting from our pipeline in dark blue. The true parameter value falls within the true posterior, but not necessarily at its centre, since

Toy-Circle (TC) model		Toy-Ellipse (TE) model	
parameters		parameters	
Name and Symbol	Range	Name and Symbol	Range
Loop radius r	[0.07, 0.23]	Loop horizontal radius a	[0.07, 0.43]
Loop center c	[0.34, 0.66]	Loop vertical radius b	[0.07, 0.43]

Table 2. The parameters of each toy model. Given parameter pair (r, c) , the TC model simulates a network whose central loop is a circle with radius $r + \varepsilon$ and center $c + \varepsilon$, where each ε is sampled from $\mathcal{U}_{[-0.03, 0.03]}(\varepsilon)$ independently and identically. The TE model uses the parameter pair (a, b) to simulate a network whose central loop is an ellipse with horizontal radius $a + \varepsilon$ and vertical radius $b + \varepsilon$ with each ε also sampled from $\mathcal{U}_{[-0.03, 0.03]}(\varepsilon)$.

the true posterior depends on $n^* = 2$ simulations of the toy model (used as observed data) which contain random noise. In all four test-cases for each model, the approximate posterior closely matches the true posterior.

To test our method of spatial model selection, we perform Step 3 of Section 3 of the main text and approximate the model posterior $p(m|\mathcal{D}^*)$ for the same four test-cases, giving results in Figure 9. Since test-cases 1 and 2 of the TC model contain circles not centred at $(0.5, 0.5)$, and test-cases 1 and 2 of the TE model contain ellipses with unequal horizontal and vertical radii, only the true model can produce their observed data. In these test-cases, we successfully infer the true model and correctly approximate posterior probabilities as 1 for the true model and 0 for the other model. The observed data in test-cases 3 and 4 of each toy model, however, could have been generated by either model. In these test-cases, the central loop in the observed data is either a circle with centre $c = (0.5, 0.5)$, or an ellipse with equal horizontal and vertical radii, which either toy model can reproduce. We therefore expect a non-zero posterior probability for each model in these test-cases, and we successfully approximate this. The approximate model posterior is still able to identify the true model in these cases, and we estimate only a small posterior probability of the incorrect mode in each test-case. While both models can simulate data similar to the observed data in each test-case, the true model will do so more often (for more parameter values). The training data from which we learn the relationship between summary statistics and model index therefore contains more simulations similar to the observed data when the true model is used, which may be why it predicts the true model with higher probability.

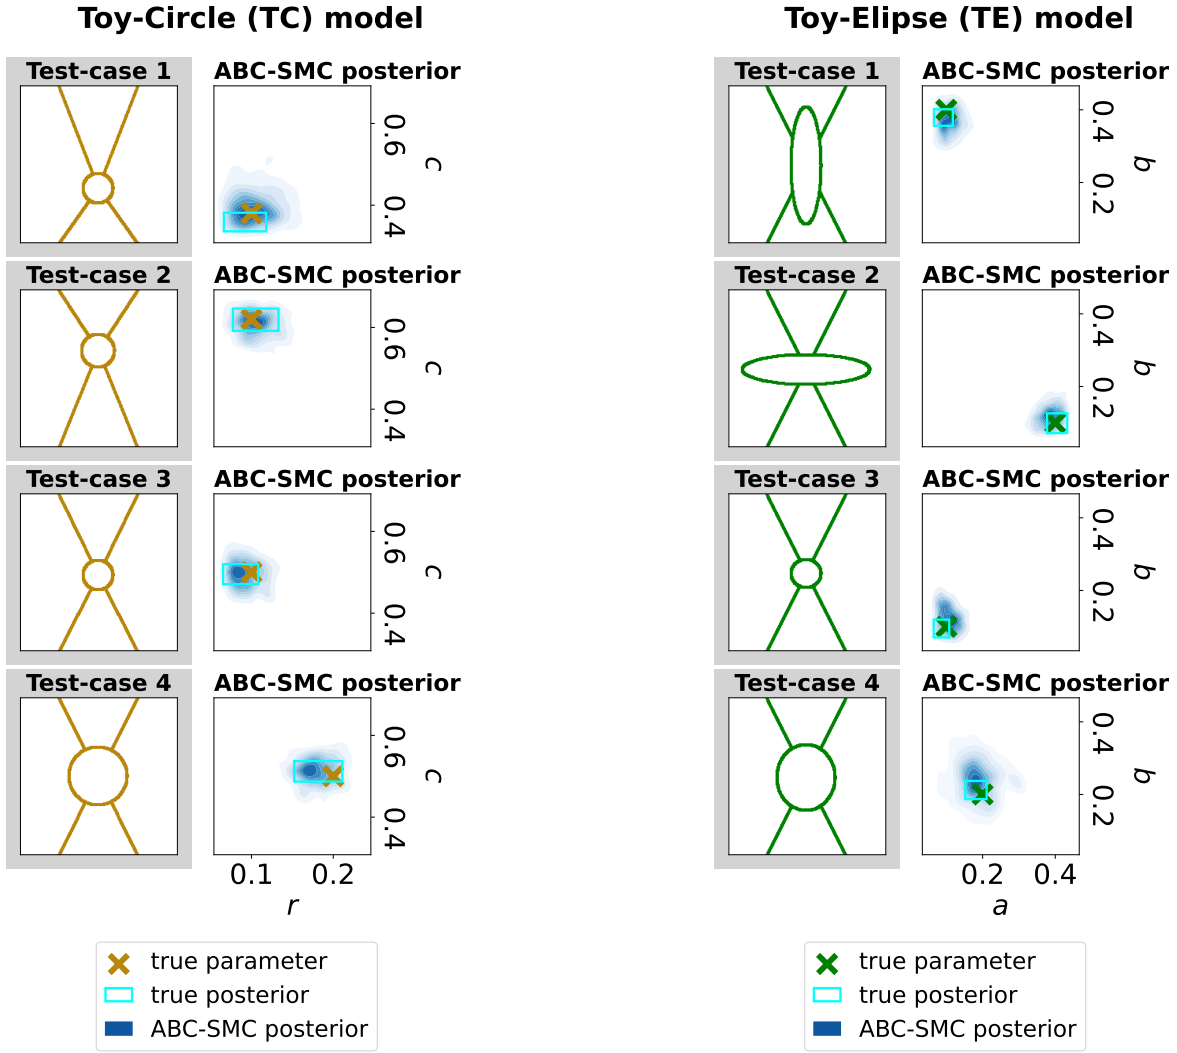


Fig. 8: Parameter inference in the TC and TE models. To generate observed data, we fix a set of parameter values and simulate the model in question $n^* = 2$ times. For each model, we pick four parameter sets to generate four test-cases and show one example of the angiogenesis datasets they simulate. We then use Steps 1-2 of Section 3 to find informative summary statistics which we use to fit each model to the observed data in each test-case. The output of ABC-SMC is a population of parameter values which approximate the parameter posterior $p(\Theta|\mathcal{D}^*)$. We plot the resulting distributions (fitting a Gaussian kernel to the parameter values accepted in the final population of the ABC-SMC algorithm). We also plot the true parameter which generated the observed data and the true parameter posterior, which we can compute exactly since the toy models are simple.

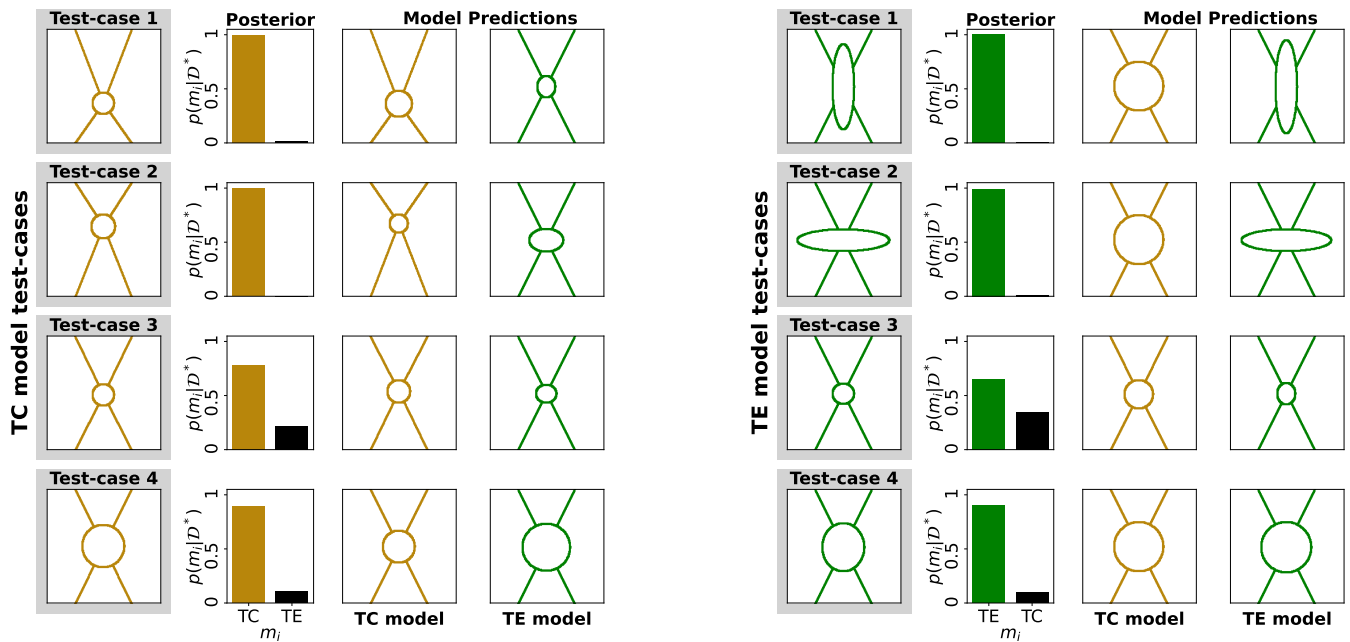


Fig. 9: Model selection among the TC and TE toy models. We use the same test-cases as in Figure 8 in which we generated observed data \mathcal{D}^* by simulating each model at known parameter values. Again we show one example of the angiogenesis datasets they simulate. Performing Step 3 of Section 3 of the main text gives an approximation of the model posterior $p(m|\mathcal{D}^*)$, which we show for each test-case. The approximate posterior selects the correct model in each test-case, giving the other model 0 posterior probability in test-cases 1 and 2 (where the observed data could only have been generated by the true model) and non-zero posterior probability in test-cases 3 and 4 (where either model could have generated the observed data). We fit each model to the observed data using Step 2 of Section 3 of the main text, and show one simulation of each model using a parameter value sampled from the approximate parameter posterior for each model. This prediction represents each model's best approximation to the observed data. In test-cases 1 and 2, only the true model predicts data similar to the observed data, whereas in test-cases 3 and 4, both models produce a visually similar approximation.

Acknowledgements

Funding

RAM thanks the EPSRC. HAH and HMB are grateful for the support provided by the UK Centre for Topological Data Analysis EPSRC grant EP/R018472/1. HAH gratefully acknowledges funding from the Royal Society RGF\EA\201074, UF150238 and EPSRC EP/Y028872/1 and EP/Z531224/1.

Data Availability Statement

All code is available at <https://github.com/rmcdmaths/tms/>. Code for the AC model was taken from Nardini *et al.* (2021) and <https://github.com/johnnardini/Angio-TDA>. The ABCpy python package <https://github.com/eth-cscs/abcpy> was used to run the ABC-SMC algorithm. Parts of Figure 1 were created in <https://BioRender.com>

References

- Adams, H., Emerson, T., Kirby, M., Neville, R., Peterson, C., Shipman, P., Chepushtanova, S., Hanson, E., Motta, F., and Ziegelmeier, L. (2017). Persistence images: a stable vector representation of persistent homology. *J. M. L. Res.*, **18**(8), 1–35.
- Ali, D., Asaad, A., Jimenez, M.-J., Nanda, V., Paluzo-Hidalgo, E., and Soriano-Trigueros, M. (2023). A survey of vectorization methods in topological data analysis. *IEEE Trans. Pattern Analysis and Mach. Intel.*, **45**(12), 14069–14080.
- Anderson, A. R. and Chaplain, M. A. (1998). Continuous and discrete mathematical models of tumor-induced angiogenesis. *Bull. Math. Bio.*, **60**(5), 857–899.
- Balding, D. and McElwain, D. L. (1985). A mathematical model of tumour-induced capillary growth. *J. Theor. Biol.*, **114**(1), 53–73.
- Bhaskar, D., Manhart, A., Milzman, J., Nardini, J. T., Storey, K. M., Topaz, C. M., and Ziegelmeier, L. (2019). Analyzing collective motion with machine learning and topology. *Chaos*, **29**.
- Blum, M. G. B., Nunes, M. A., Prangle, D., and Sisson, S. A. (2013). A comparative review of dimension reduction methods in approximate Bayesian computation. *Stat. Sci.*, **28**(2), 189 – 208.
- Breiman, L. (2001). Random forests. *Machine Learning*, **45**(1), 5–32.
- Carlsson, G. (2009). Topology and data. *Bull. Amer. Math. Soc.*, **46**(2), 255–308.

- Carlsson, G. and de Silva, V. (2010). Zigzag persistence. *Foundations of Comp. Math.*, **10**(4), 367–405.
- Chaplin, T., Harrington, H. A., and Tillmann, U. (2024). Grounded persistent path homology: A stable, topological descriptor for weighted digraphs. *Foundations of Comp. Math.*
- Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L. J., and Oudot, S. Y. (2009). Proximity of persistence modules and their diagrams. In *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry*, SCG '09, page 237–246, New York, NY, USA. Association for Computing Machinery.
- Cleveland, E., Zhu, A., Sandstede, B., and Volkening, A. (2023). Quantifying different modeling frameworks using topological data analysis: a case study with zebrafish patterns. *SIAM J. App. Dyn. Sys.*, **22**(4), 3233–3266.
- Cohen-Steiner, D., Edelsbrunner, H., and Harer, J. (2009). Extending persistence using poincaré and lefschetz duality. *Foundations of Comp. Math.*, **9**(1), 79–103.
- Connor, A. J., Nowak, R. P., Lorenzon, E., Thomas, M., Herting, F., Hoert, S., Quaiser, T., Shochat, E., Pitt-Francis, J., Cooper, J., et al. (2015). An integrated approach to quantitative modelling in angiogenesis research. *J. Roy. Soc. Interface*, **12**(110).
- Del Moral, P., Doucet, A., and Jasra, A. (2012). An adaptive sequential monte carlo method for approximate Bayesian computation. *Statistics and Computing*, **22**(5), 1009–1020.
- Edelsbrunner, H. and Harer, J. L. (2022). *Computational topology: an introduction*. American Mathematical Society.
- Farnsworth, R. H., Lackmann, M., Achen, M. G., and Stacker, S. A. (2014). Vascular remodeling in cancer. *OncoG.*, **33**(27), 3496–3505.
- Ferrara, N. (2002). VEGF and the quest for tumour angiogenesis factors. *Nature Reviews Cancer*, **2**(10), 795–803.
- Frazier, D. T., Martin, G. M., Robert, C. P., and Rousseau, J. (2018). Asymptotic properties of approximate Bayesian computation. *Biometrika*, **105**(3), 593–607.
- Hanahan, D. and Weinberg, R. A. (2011). Hallmarks of cancer: The next generation. *Cell*, **144**(5), 646–674.
- Hormuth, 2nd, D. A., Phillips, C. M., Wu, C., Lima, E. A. B. F., Lorenzo, G., Jha, P. K., Jarrett, A. M., Oden, J. T., and Yankeelov, T. E. (2021). Biologically-based mathematical modeling of tumor vasculature and angiogenesis via time-resolved imaging data. *Cancers (Basel)*, **13**(12).
- Joyce, P. and Marjoram, P. (2008). Approximately sufficient statistics and Bayesian computation. *Stat. app. in gen. and mol. biol.*, **7**(1).
- Kirk, P., Thorne, T., and Stumpf, M. P. (2013). Model selection in systems and synthetic biology. *Curr. Op. Biotech.*, **24**(4), 767–774.
- Köry, J., Narain, V., Stolz, B. J., Kaeppler, J., Markelc, B., Muschel, R. J., Maini, P. K., Pitt-Francis, J. M., and Byrne, H. M. (2024). Enhanced perfusion following exposure to radiotherapy: A theoretical investigation. *PLOS Comp. Biol.*, **20**(2), 1–32.
- Lintusaari, J., Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. (2016). Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, **66**(1), 66–82.
- Magnussen, A. L. and Mills, I. G. (2021). Vascular normalisation as the stepping stone into tumour microenvironment transformation. *Brit. J. Cancer*, **125**(3), 324–336.
- Martinson, W. D., Ninomiya, H., Byrne, H. M., and Maini, P. K. (2021). Comparative analysis of continuum angiogenesis models. *J Math Biol*, **82**(4), 21.
- Milde, F., Lauw, S., Koumoutsakos, P., and Iruela-Arispe, M. L. (2013). The mouse retina in 3d: quantification of vascular growth and remodeling. *Integr. Biol.*, **5**(12), 1426–1438.
- Nardini, J. T., Stolz, B. J., Flores, K. B., Harrington, H. A., and Byrne, H. M. (2021). Topological data analysis distinguishes parameter regimes in the Anderson-Chaplain model of angiogenesis. *PLOS Comp. Biol.*, **17**(6), 1–29.
- Otter, N., Porter, M. A., Tillmann, U., Grindrod, P., and Harrington, H. A. (2017). A roadmap for the computation of persistent homology. *EPJ Data Science*, **6**(1), 17.
- Plank, M. and Sleeman, B. (2004). Lattice and non-lattice models of tumour angiogenesis. *Bull. of Math. Biol.*, **66**(6), 1785–1819.
- Pudlo, P., Marin, J.-M., Estoup, A., Cornuet, J.-M., Gautier, M., and Robert, C. P. (2015). Reliable ABC model choice via random forests. *Bioinformatics*, **32**(6), 859–866.
- Raynal, L., Marin, J.-M., Pudlo, P., Ribatet, M., Robert, C. P., and Estoup, A. (2018). ABC random forests for Bayesian parameter inference. *Bioinformatics*, **35**(10), 1720–1728.
- Robert, C. P., Cornuet, J.-M., Marin, J.-M., and Pillai, N. S. (2011). Lack of confidence in approximate Bayesian computation model choice. *PNAS*, **108**(37), 15112–15117.
- Scianna, M., Bell, C. G., and Preziosi, L. (2013). A review of mathematical models for the formation of vascular networks. *J Theor Biol*, **333**, 174–209.
- Stepanova, D., Byrne, H. M., Maini, P. K., and Alarcón, T. (2024). Computational modeling of angiogenesis: The importance of cell rearrangements during vascular growth. *WIREs Mech. Dis.*, **16**(2).
- Stokes, C. L., Lauffenburger, D. A., and Williams, S. K. (1991). Migration of individual microvessel endothelial cells: stochastic model and parameter measurement. *J. Cell Sci.*, **99**(2), 419–430.
- Stolz, B. J., Dhesi, J., Bull, J. A., Harrington, H. A., Byrne, H. M., and Yoon, I. H. (2024). Relational persistent homology for multispecies data with application to the tumor microenvironment. *Bulletin of mathematical biology*, **86**(11), 128.
- Thorne, T., Kirk, P. D. W., and Harrington, H. A. (2022). Topological approximate Bayesian computation for parameter inference of an angiogenesis model. *Bioinformatics*, **38**(9), 2529–2535.
- Topaz, C. M., Ziegelmeier, L., and Halverson, T. (2015). Topological data analysis of biological aggregation models. *PLOS One*, **10**(5), e0126383.
- Ulmer, M., Ziegelmeier, L., and Topaz, C. M. (2019). A topological approach to selecting models of biological experiments. *PLOS One*, **14**(3), 1–18.
- Vergroesen, T. M., Vermeulen, V., and Merks, R. M. H. (2024). Falsifying computational models of angiogenesis through quantitative comparison with in vitro models. *bioRxiv:2024.08.06.606790*.
- Vilanova, G., Colominas, I., and Gomez, H. (2017). Computational modeling of tumor-induced angiogenesis. *Arch. Comp. Meth. in Eng.*, **24**(4), 1071–1102.
- Vipond, O. (2020). Multiparameter persistence landscapes. *J. M. L. Res.*, **21**(61), 1–38.
- Zomorodian, A. and Carlsson, G. (2005). Computing persistent homology. *Disc. & Comp. Geom.*, **33**(2), 249–274.